Phonetics in phonology

Chapter prepared for inclusion in the second edition of the
Blackwell's *Handbook of Phonological Theory*
(edited by John Goldsmith, Jason Riggle and Alan Yu)

D. R. Ladd, University of Edinburgh, January 2009

Final pre-publication version

1. Introduction

The primary semiotic medium of spoken language consists of acoustic signals – sound waves – produced by articulatory gestures in the human vocal tract and processed by human auditory systems. To understand more about this aspect of language it would therefore seem scientifically appropriate, even necessary, to learn more about the human vocal tract, the human auditory system, and the physics of sound. At the same time, it has been clear for more than a century that language uses the medium of sound in a very specific way, which involves the human cognitive capacity for creating categories and symbolic systems. This capacity makes it possible for two physical (acoustic) events that are objectively quite different to count as instances of the same category in the symbolic system, and for two physical events that are objectively very similar to count as instances of two different categories. It also makes it possible for different languages to categorize the physical events of speech in different ways. If we want to understand the medium of spoken language, therefore, it is not enough to consider only the physical aspects of the production, transmission and perception of sound; we need to consider the symbolic value of the sounds of speech as well.

The dual nature of speech sounds – as physical events and as elements of a symbolic system – has been recognized since the emergence of the phonemic principle in the late 19th century; in some sense, the emergence of the phonemic principle and the recognition of the dual nature of speech sounds were one and the same scientific achievement. Since the 1930s, and especially since Trubetzkoy's *Principles* (1958 [1939]), it has been customary to reserve the term *phonetics* for the study of the physical aspects of speech sounds – what Trubetzkoy described as "the study of the sounds of [Saussurean] *parole*" – and to use the newer term *phonology* for "the study of the sounds of *langue*" (Trubetzkoy 1958: 7)[1]. This terminological distinction is

[1] In citing and discussing Trubetzkoy in this paper I have carefully compared the German original with Cantineau's 1949 translation into French and with Baltaxe's 1969 translation into English, and the two translations are entered separately under the name of their translators in the list of references. In general I have taken direct quotes from Baltaxe, but the definitions of phonetics and phonology given here are my own translations of Trubetzkoy's originals *Sprechaktlautlehre* and *Sprachgebildelautlehre,* which Baltaxe renders as "the study of sound pertaining to the act of speech" and "the study of sound pertaining to the system of language" (1969: 4). I have preferred the more concise formulation in part to facilitate repeated reference to the definitions. Although Trubetzkoy's phrases do not use the Saussurean terms *langue* and *parole*, the context makes it clear that he meant to convey precisely the Saussurean dichotomy. Cantineau translates Trubetzkoy's definitions as *la science des sons de la parole* and *la science des sons de la langue* (1949: 3).

now such a fundamental part of our conceptual landscape that it seems perfectly normal for the editors of a volume on phonological theory to solicit a chapter on "phonetics in phonology". At the same time, the need for such a chapter shows that the distinction itself continues to engender conceptual difficulty. It is fairly obvious what "the sounds of *parole*" might refer to, but less obvious what "the sounds of *langue*" might be. Understanding the relation between phonetics and phonology is thus ultimately a matter of understanding the dual nature of the sign, and much of the difficulty in defining and delimiting their respective realms is ultimately due to the difficulty of deciding what sort of abstractions we are dealing with when we study language.

In the long run, the broader task of what we might call the phonetic sciences is to understand the human capacity for categorizing the sounds of speech, and to understand how this capacity reflects – and is reflected in – the structure of language. In this chapter I take some such unified ultimate goal for granted. I realize that not everyone would subscribe to it in the form in which I just stated it, and in any case there are plenty of challenging subsidiary questions on both the physical side and the symbolic side to keep researchers fully occupied without thinking about long-term goals. However, I hope to demonstrate that phonetics and phonology are inextricably intertwined even in theories that purport to draw a sharp distinction between them, and that the place of phonetics in phonology has been absolutely central ever since the phonemic principle emerged. In particular, I aim to show that many standard concepts in phonology depend crucially on the body of theory and practice that we can refer to as *systematic phonetics*. That is, most 20th century phonology – the study of the sounds of *langue* – is based firmly on a theory of phonetics – the sounds of *parole.* To the extent that there are problems with the theory of phonetics, therefore, there are problems with phonology, and this chapter also attempts to outline what some of those problems are and how we might integrate an empirically more defensible view of phonetics into our understanding of phonology.

2. Systematic phonetics in phonology

The term "systematic phonetics" is apparently due to Chomsky (1964), but the idea of systematic phonetics is embodied in the principles of the International Phonetic Association (IPA[2]). These principles are stated in summary form in successive editions of the IPA handbook, and are discussed at greater length in the handbook's most recent edition (IPA 1999) and in textbook presentations of phonetics such as Laver 1994. Systematic phonetics depends on two key premises, which I will refer to as the *segmental idealization* and the *universal categorization assumption*. These may be stated as follows:

The segmental idealization: Speech (NB not language) can appropriately be idealized as a string of ordered discrete sound segments of unspecified duration. ("Phonetic

[2] The abbreviation *IPA* is systematically ambiguous between "International Phonetic Association" and "International Phonetic Alphabet", the latter being the best-known manifestation of the former. Throughout this paper I consistently use the abbreviation only to refer to the association, and use *IPA alphabet* to refer to the alphabet. Editions of the IPA handbook are referred to here in-text as e.g. "IPA (1949)" or "IPA (1999)", because they have always been published as the work of the association, not of any specific author.

analysis is based on the crucial premise that it is possible to describe speech in terms of a sequence of segments." (IPA 1999: 5).)

The universal categorization assumption: There is a closed universal inventory of possible segment types ("The IPA is intended to be a set of symbols for representing all the possible sounds of the world's languages." (IPA 1999: 159)).

These premises were incorporated largely without comment into virtually all theorizing about phonology from the 1940s until the 1990s and are still widely accepted. Together, they yield the key theoretical construct generally known as the *phone*, and, as a kind of corollary, the notion of *(distinctive) feature.* These ideas are now such a fundamental part of the way we think about phonetics that it comes as a surprise to realize that they were not taken for granted until well into the 20th century, and it is worth taking some time to trace their development.

2.1. The phone

The phone has been part of the IPA enterprise from the very beginning, but at first it was only implicit. According to the history of the International Phonetic Association included in the IPA Handbook (IPA 1999: 194-197), the IPA started out life in 1886 as, in effect, a response to the inconsistencies of English orthography, aiming at a practical orthography with consistent phoneme-grapheme correspondences for use in language teaching. However, the idea of developing a consistent practical orthography adaptable to all languages was explored very early in the history of the IPA, and the first version of the IPA alphabet was published in 1888, along with a set of principles on which it was based. The first of these principles (again, according to IPA 1999) was:

> "There should be a separate sign for each distinctive sound; that is, for each sound which, being used instead of another, in the same language, can change the meaning of a word".

In modern terms, this very clearly states that IPA transcription is intended as a phonemic transcription, and *sound* is clearly being used to mean "phoneme". However, the seeds of theoretical confusion were sown immediately, in the second principle:

> "When any sound is found in several languages, the same sign should be used in all. This applies also to very similar shades of sound."

This second principle requires us to define "sound" in a different way from the first principle, because we cannot use any sort of practical test based on word meaning to decide whether two sounds in two different languages are distinctive or not. The notion of sound in the first principle is language-specific; the notion of sound in the second implies a language-independent categorisation. This second sense of sound is what came to be known as the phone.

Leonard Bloomfield, the central figure in American linguistics in the first half of the 20th century, saw the contradiction between these two principles and devoted several pages of his great work *Language* (1933) to trying to expose it. He uses the term

"phonetic [sic] transcription" to refer to "a system of written symbols which provides one sign for each phoneme of the language we are recording", and explicitly denies the validity of attempts to transcribe non-distinctive acoustic detail. It is worth quoting him at some length:

> Having learned to discriminate many kinds of sounds, the phonetician may turn to some language, new or familiar, and insist upon recording all the distinctions he has learned to discriminate, even when in this language they are non-distinctive and have no bearing whatever. … The chief objection to this procedure is its inconsistency. The phonetician's equipment is personal and accidental; he hears those acoustic features which are discriminated in the languages he has observed. Even his most 'exact' record is bound to ignore innumerable non-distinctive features of sound; the ones that appear in it are selected by accidental and personal factors. … [H]is most elaborate account cannot remotely approach the value of a mechanical record.
>
> Only two kinds of linguistic records are scientifically relevant. One is a mechanical record of the gross acoustic features, such as is produced in the phonetics laboratory. The other is a record in terms of phonemes, ignoring all features that are not distinctive in the language. … (pp. 84-85)

However, Bloomfield's views had essentially no influence on subsequent theoretical developments, not even among his closest followers, the so-called neo-Bloomfieldians like Bernard Bloch (e.g. 1941, 1948) and Charles Hockett (e.g. 1942, 1955). Instead, the idea that there is a valid universal basis for abstracting segment-sized sounds out of the stream of speech, and a valid universal framework for categorizing them, became firmly established in the 1920s and 1930s.

It is true that there was at least one attempt to put the phone idealization on a firm theoretical footing. In his 1943 monograph *Phonetics*, Kenneth Pike devoted an entire chapter (entitled "Units of sound") to the theoretical difficulties with the notion "speech sound" or "phone", stating the problem as follows:

> "Speech, as phoneticians well agree, consists of continuous streams of sound within breath groups; neither sounds nor words are separated consistently from one another by pauses, but have to be abstracted from the continuum. Phonemicists concur in the belief that some unit of speech, the phoneme, can be discovered as the basic constituent of a linguistic system. … Is there a significant halfway point between the continuum and the phoneme? Is there a real, nonfictitious segment of sound which is not a phonemic one?" (p. 42)

Bloomfield's answer to Pike's question, as we just saw, was unambiguously "No." Pike, however, after some discussion of differing views and difficult cases, inclines toward "the conclusion that there must ultimately be some such phonetic segmentation behind speech" (p. 46). He then sets out (p. 52) to find "a workable method for the delineation of natural phonetic segmentation", in which the "segmental unit is to be determined entirely apart from phonemic function". He notes that "[a] corollary of this aim states that such a segmentation procedure is equally applicable to any and all languages, or to any stream of nonsense syllables." Such a procedure means that "an impressionistic phonetic record of a new language proves

theoretically legitimate as well as practically valuable … for the phonemicist …" (p. 53, emphasis added).

It is difficult to know to what extent Pike's theoretical considerations influenced the development of the field, but it is clear that few writers after him were worried about the theoretical legitimacy of the phone idealization, or about the assumption that there is a closed universal set of phones. By 1949 (when the IPA *Principles* were republished in revised form), the notion of discrete speech sounds or phones appears to be taken for granted. The new version of the first principle starts: "When two sounds occurring in a given language are employed for distinguishing one word from another..."; and the second begins: "When two sounds are so near together acoustically that there is no likelihood of their being employed in any language for distinguishing words..." (IPA 1949: 1).

Pike's reference to the practical value of his procedures for the "phonemicist" reminds us of the central role that systematic phonetics had already come to play in theoretical phonology. During the 1930s and 1940s the phone idealization became firmly embedded in linguistic discussions of the phoneme on both sides of the Atlantic – as for example in Trubetzkoy's discussion of how to define and identify phonemes (1958: chapter II), which simply presupposes the phone (*Lautgebilde,* translated as *sound* by Baltaxe; cf. her translator's note (1969: 36)). Early codifications of the "allophone" idea (e.g. Bloch 1941, Hockett 1942; cf. Trubetzkoy's "combinatory variant") are probably the clearest illustration of the central importance of the phone concept in shaping phonological theory.

Consider the realisation of voiceless stops in English syllable onsets, which is probably used as an example in 90% of beginning linguistics courses in the English-speaking world. It is well known that in absolute initial position, as in *peach*, voiceless stops typically have a voice onset time (VOT) in the general range of 50-70 ms, whereas when preceded by /s/ in an onset cluster, as in *speech*, they typically have a VOT in the general range of 0-20 ms. This is an easily observable fact about the phonology of English, and provides a clear and simple illustration of the fundamental phonological concept of lawful conditioned variation. However, statements of this variation are conventionally expressed not in terms of mean VOT, but in terms of two phones, usually notated e.g. [p] and [p^h], the latter occurring in absolute initial position and the former occurring after /s/. This statement is already a considerable abstraction away from observations about VOT, but that is not acknowledged in most classical formulations of the phoneme or in most textbook presentations. Instead, the phones are considered to be the *raw data*; transcriptions like [spitʃ] and [p^hitʃ] are assumed to provide a faithful representation of what a speaker really produces. Rather than recognise [p] and [p^h] as abstractions based (as Bloomfield emphasized) on the personal equipment of the transcriber[3], classical

[3] Actually, it is probably not quite accurate to say that phone abstractions are based on the personal equipment of individual phoneticians; it would be better to describe them as the collective effect of the personal equipment of a group of scholars who were all literate in alphabetic writing systems and all spoke more than one European language. Alphabetic literacy inclined them toward the segmental idealization; familiarity with several languages that used the same alphabet inclined them to identify cross-linguistic categories of sound like [b] and [p] and to focus their attention on specific phonetic details (like the difference between [p] and [ph]) that were salient in the comparative description of the European languages.

phoneme theory took them as categories of phonetic description, identifiable in a language-independent way.

I am, of course, well aware of operating with the benefit of hindsight here. When I say that the facts about English VOT are "easily observable", I am referring to the technological environment of today, not that of the 1950s or even the 1980s. Today, given free software like Praat (http://www.praat.org) and readily accessible tools for plotting data, it is indeed a simple matter to establish that the facts of English VOT are roughly as I have stated them and to see clearly that such facts are a matter of statistical distributions, not unvarying categories. However, an early attempt to base phonology on such a statistical view of phonetics (Zwirner's "Phonometrie"; Zwirner and Zwirner 1966 [Zwirner 1936]) was rejected by Trubetzkoy (1958: 10-12) in what may be seen as an early instance of the gulf of misunderstanding between phoneticians and phonologists. Even after the sound spectrograph brought general awareness of the variability of the raw data – triggering considerable soul-searching at least on the part of neo-Bloomfieldian phonologists (see e.g. Joos 1948; Bloch 1948, esp. footnote 6 and postulates 9 and 11; Hockett 1955, section 5) – the phone idealization always managed to survive[4].

The supposed reality of phones was crucial to the role played in traditional definitions of the phoneme by the minimal pair test, i.e. the substitution of one sound for another. Postulating a phonemic distinction between /p/ and /b/ in English depends in part on agreeing in advance that [pʰ], [p] and [b] are comparable sounds or segments in pairs like *pit/bit, pang/bang, cap/cab, poor/boor,* and so on. In the case of [pʰ], [p] and [b], there is little disagreement that these are comparable units, but there are many well-known cases where there was no such agreement and phonemic analysis was correspondingly controversial. The best-known case in English is probably that of the affricates, and the problem of whether to treat affricates and other such complex segments as single phonemes or as clusters has a long history. The relevance of segmentation to these cases is as follows: if *chip* begins with the phones [t] and [ʃ], then [t] can be replaced by zero to yield *ship* and [ʃ] by [ɹ] to yield *trip*, so that *chip* can be said to begin with a cluster; if, on the other hand, we do not identify the first part of the affricate with the phone [t] and/or do not identify the second part with [ʃ], then there is no obstacle to treating the affricate as one phone and analysing the beginning of *chip* as a single consonant. Without a universally valid method of segmentation *and a universally valid system of classifying segments as the same or different*, defining phonemes in terms of the distribution of phones is ultimately arbitrary, as Pike correctly saw. Pike's faith that such a segmentation could be justified theoretically was not shared by e.g. Martinet (1966 [1939]), who says: "From all this, it turns out that the first task of the phonologist is an in-depth phonetic analysis of the language under study, during which analysis it will be necessary above all to be careful not to be led astray by the imperfections of traditional phonetic transcriptions" (p. 122, my translation). In other words, Martinet recognizes that the identification of the phones on which we base our theoretical definition of the phoneme is specific to a given language.

[4] The same soul-searching still goes on among self-identified phoneticians thoroughly familiar with the continuous parametric nature of speech. A particularly striking example is seen in Laver's defense of systematic phonetics (1994, section 4.4), which comes close to acknowledging that a symbolic segmental representation cannot be reconciled with what we know from instrumental research.

Nevertheless, 20th century theories of phonology were universally built on the assumption that phones and phonetic transcriptions are a scientifically appropriate language-independent representation of speech. This was the idea that Chomsky picked up in his brilliant dissection (1964) of mid-century phoneme theory and his presentation of the assumptions underlying what became mainstream generative phonology. He drew a sharp distinction between "physical phonetics" and "systematic phonetics", explicitly claiming that both levels of description are necessary in a formal model of language and speech. Specifically, he envisaged an overall theoretical structure in which the output of the phonology (or, more broadly, the output of the grammar) is a systematic phonetic representation consisting primarily of a string of phones; this systematic phonetic representation is then passed to a phonetic implementation system – not part of *langue* – where universal biomechanical and physical principles generate the physical phonetic output. In terms that have become familiar more recently, generative phonology thus sees the systematic phonetic representation as the *interface* between phonology and phonetics – or, if we accept Trubetzkoy's definitions, the boundary between *langue* and *parole*. As is well known, Chomsky argued that the "taxonomic phonemic" level of the neo-Bloomfieldians was unnecessary and unmotivated, and that the phonological grammar should map directly from abstract "systematic phonemic" representations to the systematic phonetic output (cf. also Halle 1959). Like the neo-Bloomfieldians, however, he did not question the assumption that the systematic phonetic representation is a scientifically valid idealization. Indeed, this assumption was vigorously defended by Postal (1968) and with very little further discussion was incorporated into the generative theory codified in *SPE* (Chomsky and Halle 1968) and a number of textbook presentations in the 1970s (e.g. Schane 1973, Hyman 1975, Kenstowicz and Kisseberth 1979).

Since the 1960s, few phonologists have questioned the early generative acceptance of systematic phonetics and the segmental idealization, and the idea of universal phonetic categorization remains at the foundation of most present-day work in phonology. It is true that in the late 1980s and early 1990s there was a flurry of interest in interface issues. In 1990 the question of phonetic representation occupied an entire special issue of the *Journal of Phonetics* (vol. 18: 297-477), in which the segmental idealization was attacked (e.g. Pierrehumbert 1990), assumed (e.g. Ladefoged 1990), and defended with empirical evidence (e.g. Nearey 1990). However, at more or less the same time the attention of the field was captured by Optimality Theory (OT; e.g. Prince and Smolensky 2004, Archangeli and Langendoen 1997, Kager 1999) and interface issues were largely marginalized. OT incorporates the generative understanding of phonetics wholesale: its entire architecture is based on having a set of categorically distinct "outputs" to evaluate, which is possible only if we abstract away from the infinite variability of speech and assume some sort of universal categorization of the speech sounds. Moreover, the key faithfulness constraints with which the theory began, PARSE and FILL (and their successors MAX and DEP), are built on the assumption that the output can be exhaustively and unambiguously divided into segments. Within OT, there have been some attempts to deal with the empirical difficulties posed by these assumptions (notably Boersma 1998), but the great body of work in OT continues to accept systematic phonetics as a valid basis for describing the output of the grammar, and as a convenient delineation of the boundary between its concerns and those of others.

2.2. Distinctive features

The idea of a universal scheme of classification for phones gives rise to what is perhaps the central theoretical construct of mid-20th century phonology, namely the feature. In an informal way, of course, the dimensions of the IPA symbol chart are a kind of feature analysis, but we are concerned here with the place of such classification *in phonology*. Linguists had long been aware that certain kinds of sound changes are common and somehow natural, and that common phoneme inventories across languages are often quite symmetrical if described in terms of phonetic dimensions. But this awareness played no formal role in most Anglo-American phonemic theorizing, which was almost exclusively concerned with the procedures for grouping phones into phonemes. The work of putting phonetic symmetries and similarities on an explicitly phonological footing was carried out by the members of the Prague School during the 1930s.

The basic ideas were presented by Trubetzkoy in *Principles*. Trubetzkoy's theoretical starting points were, first, the strict separation of phonetics and phonology, and second, the structuralist or Saussurean idea that language involves a system of oppositions, in which the central property of any given sign is that it is *not* any of the other signs. This last idea is the view summed up in Saussure's well-known dictum "Dans la langue il n'y a que des différences" and in Jakobson and Halle's suggestion (1956: 22) that the meaning of a phoneme is "mere otherness". Accordingly, Trubetzkoy starts out by describing phonology in purely abstract terms ("The signifier of the system of language [i.e. of *langue*] consists of a number of elements [viz., phonemes - DRL], *whose essential function it is to distinguish themselves from each other.*" (Baltaxe 1969:10, emphasis added)). Nevertheless, in order to talk about the actual systematic differences that distinguish one phoneme from another – differences in *langue* – Trubetzkoy did not refer to abstract dimensions but to concrete phonetic properties of phones – elements of *parole*. He treats this recourse to phonetic dimensions as inevitable: "As regards phonology, it is clear that it must make use of certain phonetic concepts. For instance, the claim that in Russian the contrast between voiced and voiceless obstruents is used in Russian to differentiate between words belongs to the field of phonology. The terms 'voiced' and 'voiceless' and 'obstruents' themselves, however, are actually phonetic" (1969: 14). He reiterates the necessity of making this link to phonetic concepts at greater length in the introduction to Chapter IV (1969: 91f).

Trubetzkoy's version of features (or "oppositions") was thus in some important respects merely an expedient codification of the dimensions of the IPA chart. In particular, the distinction he draws among "privative", "gradual" and "equipollent" oppositions is patently related to – if not actually influenced by – the physical nature of those dimensions, and much of his discussion is cast in traditional IPA terms. However, three major subsequent developments meant that the feature concept took on a theoretical life of its own.

The first development was the publication of Jakobson, Fant and Halle's *Preliminaries to Speech Analysis* (1952; henceforth *JFH*), which presented a fully worked out theory of distinctive features whose dimensions were no longer merely those of the IPA. The most conspicuous taxonomic innovations were that the features were exclusively binary and that they were defined in purely acoustic terms.

However, the *JFH* feature system reaffirms the two premises of systematic phonetics identified at the beginning of this discussion: it presupposes the segment, and it explicitly presents the taxonomic framework as universally valid. Actually, the *JFH* version of the segmental idealization does represent a refinement of the IPA version, because it acknowledges the continuous variation of acoustic parameters. Specifically, it treats the segment not as a section of signal with duration, but as an idealized instantaneous slice through the signal at a specific point in time: "For practical purposes each phoneme can be represented by a quasi-stationary spectrum in which the transfer function is invariable with respect to time…." (1952:18). The features that characterize the segment are therefore based on the acoustic properties of the signal at the point in time when the idealized instantaneous slice is taken. Apart from that refinement, however, the *JFH* approach is built on a conception of the phone that was perfectly consistent with the ideas of Pike or Hockett.

The *JFH* definition of segment, and the concomitant definition of the feature as an actual acoustic property at an identifiable point in time, is part of a second important development in feature theory that is much less widely recognized. For Trubetzkoy, features are above all *abstract* characteristics of phonemes: phonemes are the elements of phonology, forming part of a system of *oppositions*, and phonetic properties are of interest only insofar as they describe how the abstract oppositions are manifested. The proposal in *JFH* that phonemes are instantaneous time slices at which features can be identified *in the signal* represents a considerable departure, in that the features have become acoustic events or properties of acoustic events rather than abstract dimensions. This in turn easily leads to the idea that the elements of phonology are features, and phonemes are composite. Such a conception is strongly suggested by *JFH* and made explicit by Chomsky and Halle's work in the 1950s and 1960s, but is clearly absent from Trubetzkoy's thinking.[5]

This finally brings us to the third major development of the feature notion, namely its incorporation into the phonological theory of *SPE*. In some respects the *SPE* version of feature theory was conservative: it did not question the assumption that features should provide a universal framework for describing actual sounds, and it did not pursue the *JFH* definition of the segment as an instantaneous time-slice, conceiving of

[5] When Trubetzkoy discusses the phonetic basis of oppositions he normally uses the German word *Eigenschaft*, which is quite abstract and is appropriately translated into English as *characteristic* or *property*; he seldom uses the word *Merkmal*, which is now the standard German technical term for the modern sense of 'feature', and which more clearly conveys the idea of an actual mark of some sort. The English word *feature* is much more ambiguous: it can refer not only to abstract characteristics but also to specific objects or actual marks of some sort, especially in fixed collocations like 'features of the landscape' or 'distinguishing feature (of a person)'. Cantineau generally translates *Eigenschaft* as *particularité* or *caractéristique* rather than *trait,* which is now the standard French technical expression for the modern sense of 'feature'; when *Merkmal* occurs Cantineau generally renders it as *marque*. Baltaxe, who prepared her translation in the late 1960s after the technical use of *feature* was well-established, deliberately avoided the term *distinctive feature* and carefully distinguishes "Trubetzkoy's theory of distinctive oppositions" from "[Jakobson's] theory of 'distinctive features'" (1969:vi-vii); she generally renders *Eigenschaft* as *property* and *Merkmal* as *mark*. To the extent that one can carry out a non-electronic search of a text as long and as dense as *Principles*, it appears that the only place Cantineau uses the phrase *trait pertinent* is at the beginning of chapter III, where Trubetzkoy (1958: 59) describes the "phonological content" of a phoneme as the *Inbegriff aller phonologisch relevanten Eigenschaften* (NB not *Merkmale*), which is translated as "all phonologically distinctive properties" by Baltaxe (1969: 66) and as "l'ensemble des traits phonologiquement pertinents" by Cantineau (1949: 68).

sounds very traditionally as phones. However, it formally adopted the notion that features are the primitive elements of phonology and phonemes merely sets or "bundles" of such primitives. Moreover, it took seriously another idea, implicit in Trubetzkoy but not developed in *JFH*, namely that the universal descriptive framework established by the set of features should also allow us to express phonological symmetries and generalizations. This led to the better-known aspect of Chomsky and Halle's revision of *JFH*, namely the replacement of several of the acoustically-based *JFH* features such as [grave] and [compact] by features based on articulatory dimensions more like the traditional dimensions of the IPA chart. The principal justification for these changes was that the new features were better suited to expressing the generalizations of phonology. Like Trubetzkoy, that is, Chomsky and Halle seem to have concluded that the best way to give a description of phonological regularities was in terms of the taxonomic dimensions of phonetics.

Considering the importance that Trubetzkoy attached to the phonology-phonetics distinction, the persistence of traditional phonetic dimensions in phonology is striking. One could perfectly well imagine a description of the distinctive oppositions in a given language that makes no reference to phonetics and really does work with the idea of abstract distinctness or "mere otherness". Standard names for the four tonemes of Mandarin Chinese are essentially of this sort: the long-standing Western practice of using the numbers 1 to 4 obviously makes no reference to the pitch contours by which the abstract tonemes are phonetically manifested. (Essentially the same is now true of the traditional Chinese names *yīn píng* 'yin level', *yáng píng* 'yang level', *shàng* 'upper', *qù* 'leaving', though in Classical Chinese these may have had some phonetic content.) Indeed, this might seem to be a good way of pursuing Trubetzkoy's professed goal of categorizing "the sounds of *langue*": such names or numbers are shorthand ways of referring to abstract phonological elements that are functionally equivalent across the lexicon irrespective of phonetic realization. For example, "Tone 2" is mid-high-rising in standard Mandarin and mid-low-falling in Chengdu (Chang 1958). The phonetic realization could hardly be more different, but the *system* of tones in both varieties is still basically the Mandarin four-tone system, in the sense that words having "Tone 2" in one variety will reliably have it in the other as well.

It is true that non-phonetic names like "Tone 2" are names for whole phonemes, not features, but there is no obvious reason why non-phonetic names could not also be used to designate the patterns of opposition that Trubetzkoy saw as the essence of phonology. Indeed, it is not hard to see that phonetically abstract names for phonologically relevant dimensions are sometimes exactly what we want. Perhaps the clearest instance is Chomsky and Halle's proposal for a feature [syllabic] to replace the *JFH* feature [vocalic]. Although they provide an ostensibly phonetic definition of [syllabic] as "constituting a syllable peak" (1968: 354), they give little indication of the difficulty of defining syllables phonetically, and the motivations for having such a feature are patently phonological. Similar remarks could be made about the feature [tense] applied to vowels in English or Dutch, or about the descriptive term *rhotic,* which is sometime used to refer to the phonetically diverse set of segment types that manifest the /r/ phoneme in English and other European languages.

Nevertheless, the unquestionable descriptive utility of such phonetically abstract features has not so far raised any serious theoretical doubts about the appropriateness

of using phonetic dimensions to characterize phonological oppositions. On the contrary, a good deal of theoretical work (e.g. Hayes and Steriade 2004) has examined the "grounding" of phonological features in phonetics, and the phonetic basis of feature definitions is now seen as involving a significant theoretical claim, "namely, that natural phonological classes and sound changes will be definable in phonetic terms" (Kenstowicz and Kisseberth 1979: 240). Following Postal 1968, Kenstowicz and Kisseberth refer to this claim as the "naturalness condition" and assume its validity. For example, they say explicitly of the feature [syllabic] that "[s]ince the syllable has not yet been defined satisfactorily in phonetic terms, the phonetic correlates of this feature are unclear" (1979: 242), implicitly presupposing that such satisfactory phonetic definition will eventually be forthcoming. This presupposition is made explicit when they note more generally that "there are still a number of widespread phonological processes which presuppose natural classes of sounds for which no straightforward phonetic correlates are presently known. They pose a challenge to future research and one can only hope that as phonetic science progresses, these unexplained counterexamples to the naturalness condition will eventually be resolved" (1979: 241). In short, they treat any difficulties in reconciling phonetic and phonological uses of features as a matter for empirical research rather than theoretical reconsideration.

3. Systematic phonetics in its own right

In the discussion so far I have sought to show that a crucial component of most contemporary conceptions of phonology is a theory *of phonetics*: the rigid separation between phonetics and phonology posited by Trubetzkoy and assumed by subsequent generations of linguists is illusory (cf. also Chomsky 1964: 109f [1972: 423]). The illusion could be harmless, of course. As long as the theory of phonetics is approximately valid, then what I have said so far amounts to little more than an academic exercise in the exegesis of classic texts. That is, it could be that Trubetzkoy was wrong about the strict division, but nothing else of substance changes. In this case, Kenstowicz and Kisseberth would be justified in awaiting the results of further empirical progress in phonetic science.

However, since the 1980s progress in phonetic science has been considerable. The increasing ease of acquiring instrumental data – especially acoustic data, but also articulatory data – means that we know more and more about the details of phonetic realization. Much of this research has been carried out under the heading of "laboratory phonology" (e.g. Kingston and Beckman 1990, Pierrehumbert et al. 2000), a phrase that would probably have struck Trubetzkoy as an oxymoron. But the phrase is precise and meaningful: laboratory phonology examines the sounds of *parole* not in order to learn more about the processes of speech production and perception, but to evaluate the implicit predictions that phonological representations make about phonetic behavior (cf. the discussion in Beckman and Kingston 1990). Little of what has been found is compatible with the phonetic idealizations that – as we have seen in the foregoing sections – underlie modern phonology. Indeed, there is now plenty of reason to think that there are serious problems with systematic phonetics as a theory of speech. These problems are briefly sketched here.

3.1. Systematic phonetics as universal categorization

The first set of problems with systematic phonetics involves the goal of providing a universally valid taxonomy of speech sounds. Even before the advent of cheap and accessible acoustic analysis, some traditional phoneticians commented on the Eurocentric bias in the IPA's categories, but recent instrumental work makes it increasingly difficult to maintain the idea of a universal categorical taxonomy. A striking example comes from Cho and Ladefoged's careful comparative study (1999) of voice onset time (VOT) in eighteen different languages. Figure 1 shows the mean VOT in voiceless velar stops in citation forms before non-high vowels for each of the languages; in some cases the languages in question had two such phonemes, one with short-lag ("unaspirated") and one with long-lag ("aspirated") VOT. It can be seen that there is a more or less continuous range of mean VOT values; there is certainly nothing like a cluster for unaspirated and a cluster for aspirated. The authors do suggest that the continuum might be divided up into four regions (indicated by the boxes in Figure 1) called "unaspirated", "slightly aspirated", "aspirated" and "highly aspirated", but this view strikes me as implausible, especially considering the relatively small size of the sample of languages. That is, it seems very likely that if we computed means from many more languages with the same methodological rigor, any apparent discontinuities in the gradual increase from one end of the VOT scale to the other would disappear.

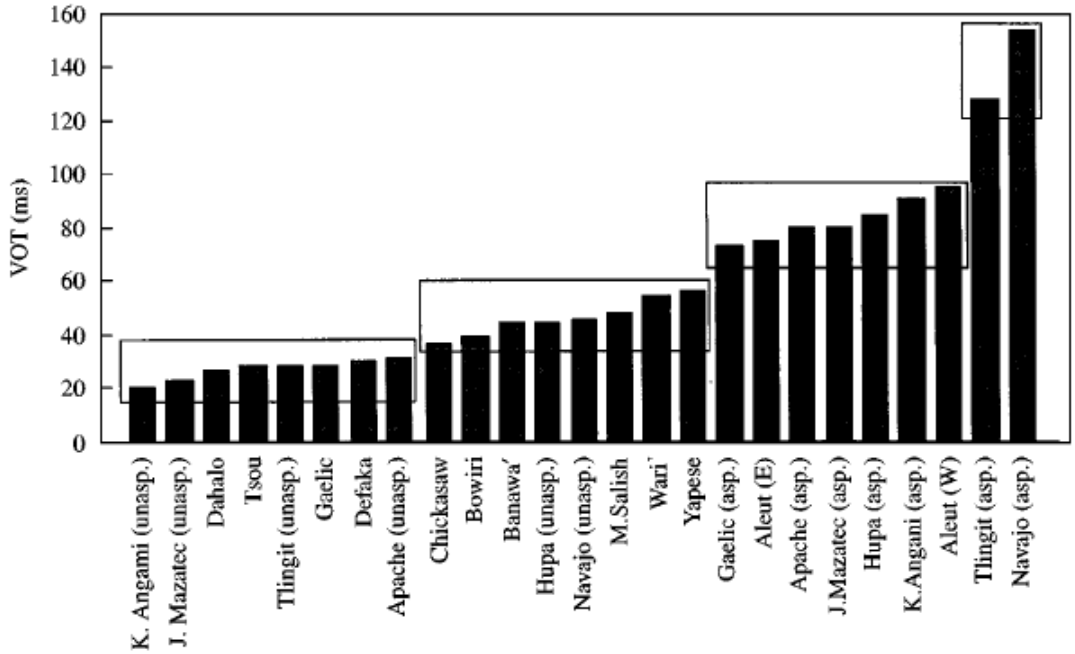

Figure 1. Mean voice onset time for 25 voiceless stop phonemes in 18 languages. From Cho and Ladefoged 1999.

A different kind of challenge to any notion of universal categorization comes from recent work on Kera (a Chadic language spoken by some 50,000 people in Chad) by Mary Pearce (2007). According to a standard phonetic and phonological description (Ebert 1975-79, taken up by Odden 1994 and Rose and Walker 2004, all cited in Pearce 2007), Kera has both voiced and voiceless stops and three distinctive tones

(high, mid and low), with various co-occurrence restrictions (in particular, voiced stops occur primarily before low tone). By analysing the productions of several Kera speakers acoustically, however, Pearce showed that in fact VOT is extremely variable in all stops, and *co-varies with pitch*: as shown in Figure 2, VOT has the shortest mean in low toned syllables and is slightly longer in mid and high toned syllables, but the VOT ranges of all three tones substantially overlap. That is, VOT is not distinctive in Kera, but some of the variation in VOT is predictable from tone, and therefore, in effect, VOT is one of the phonetic cues to tone. The two-way categorization of stops as voiced or voiceless is based on the Eurocentric categories of the first phoneticians to describe the language – exactly the kind of thing Bloomfield warned against in the passage quoted earlier. Moreover, the idea that VOT could serve as a phonetic cue to the phonological category of tone cuts across a standard understanding of the distinction between segmental and suprasegmental.[6] But as Pearce amply shows, the description of the phonology of Kera makes much more sense if we adopt exactly that idea.

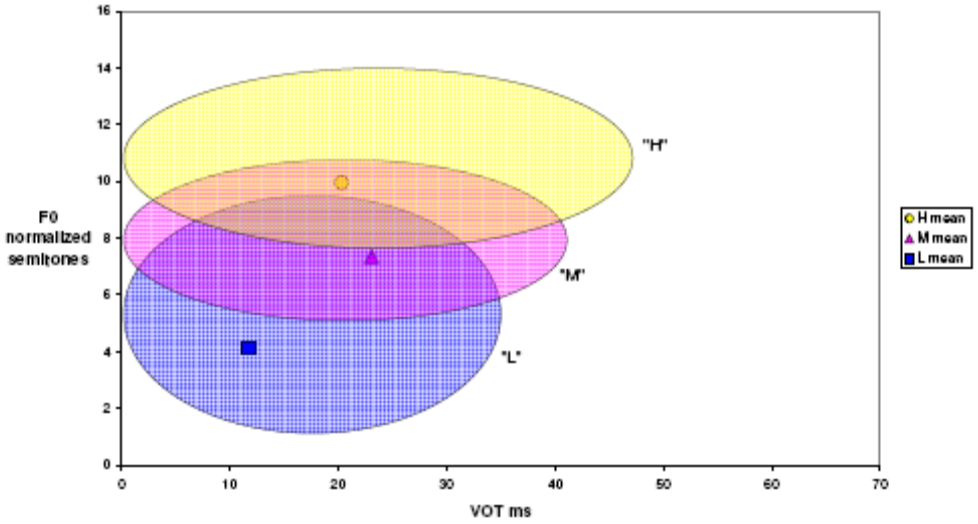

Figure 2. Summary plots of voice onset time and fundamental frequency for Kera syllables with high, mid, and low phonological tone. This figure is based on village women's speech; village men's speech is similar. Town-dwelling speakers influenced by French show slightly clearer VOT-based distinction between low and the other two tones. From Pearce 2007.

Both the cases just discussed could be incorporated into a modified systematic phonetic theory in which phones are defined in language-specific quantitative terms as a mean value on some measurable phonetic scale (or, more generally, as a central value in some quantitatively definable phonetic space such as the vowel space defined by the value of the first two formants). That is, we could give up the idea of *universal* categorization, but still maintain the segmental idealization and still maintain the idea that the output of the phonology is a string of systematic phones which are then passed on to physical phonetics for realization. Such realizations could be quite variable without upsetting the quantitative definition of the phone. As noted above

[6] The distinction between segmental and suprasegmental is arguably another consequence of systematic phonetics, "suprasegmental" properties being merely those that are left over when an utterance is divided into phones. This topic is discussed briefly in section 3.3 below.

(in the discussion of VOT allophony in English) and as just illustrated in Figure 2, it is now quite normal to describe the phonetic manifestation of a given phoneme in statistical terms: specifically, it is common to present such data graphically as a distribution ("cloud") of individual realization tokens in some appropriate phonetic space, and it is normal to find that the edges of such clouds overlap, even quite considerably. None of this need threaten the idea that language-specific allophones can be defined quantitatively, each with its own portion of phonetic space, as long as the overlapping distributions are statistically distinct.

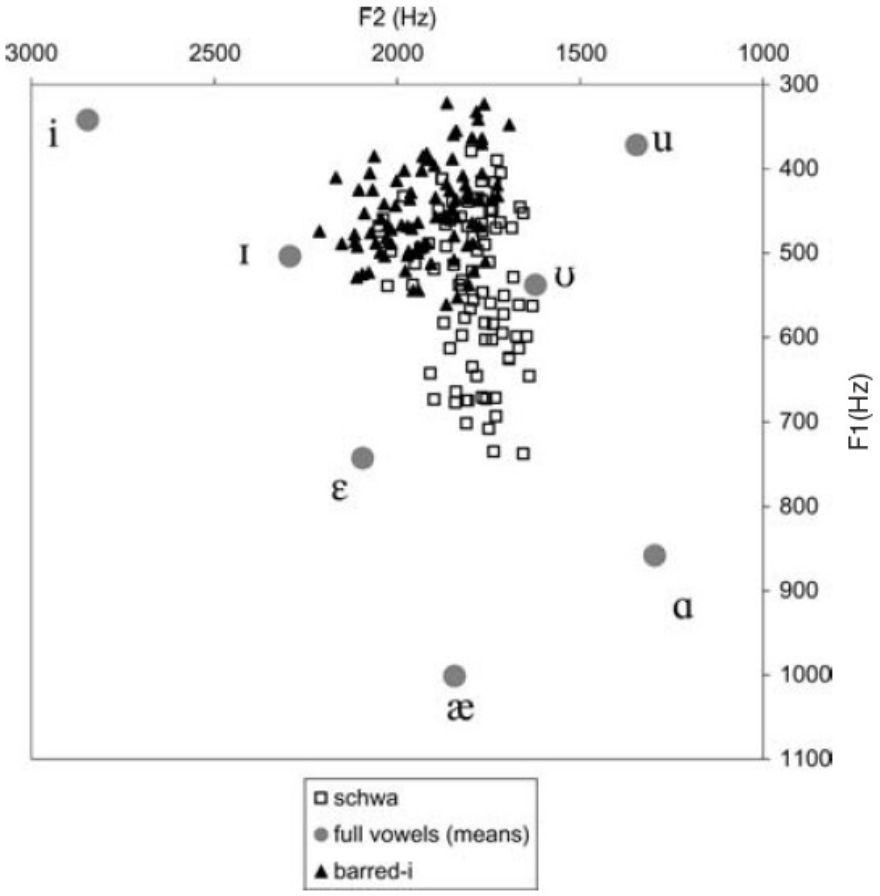

Figure 3. Scatterplot of formant values of individual tokens of phonemic /ə/ (from Rosa's) and phonemic /ɨ/ (from roses). From Flemming and Johnson 2007.

However, even this idea is hard to reconcile with the results of another recent study. Flemming and Johnson (2007) investigated the acoustic realization of the two unstressed vowels found in phrases like *Rosa's roses* in American English. The two vowels are clearly distinct, in the sense that formant plots of multiple tokens of each vowel show different distributions: the second vowel of *roses* is on average higher than that of *Rosa's*, which seems to justify transcribing the two with, say, [ɨ] and [ə] respectively, as Flemming and Johnson suggest. However, the way in which the distributions overlap, shown in Figure 3, means that [ɨ] is essentially a subset of [ə]. There is no obvious way to reconcile this kind of distributional fact with a traditional phone-based transcription. A traditional approach might be to say that there are two distinct phones [ɨ] and [ə], one of which is used in *roses* and either of which can be used "in free variation" in *Rosa's,* and careful IPA transcriptions might represent the

greater variability of *Rosa's* in exactly that way. But it can be seen that this description misrepresents the quantitative data: the distribution of the vowel in *Rosa's* appears to occupy a continuous space on the plot, not two separate spaces corresponding to two different transcriptions. That is, the quantitative data justify the statement that there are two distinct unstressed *phonemes* /ɨ/ and /ə/ in American English, but not that American English phonetic realizations allow us to distinguish two *phones* [ɨ] and [ə] occupying reasonably distinct areas of phonetic space.

3.2. Systematic phonetics as interface representation

The second set of problems with systematic phonetics revolves around the notion of *interface*. As we have already noted, systematic phonetics is often seen, even by scholars of very different persuasions, as a *level of representation* at the *interface between the abstract and the physical.* This understanding of systematic phonetics is made explicit in generative phonology, beginning with Chomsky 1964 and Postal 1968, but it is implicit, as Chomsky saw, in the IPA idea that there is a universally valid segmental representation of *utterances* in any language. Such an understanding is what lies behind Pike's question "Is there a significant halfway point between the continuum and the phoneme?" Some of the discussions of IPA transcription in the 1990 special issue of *Journal of Phonetics* mentioned earlier focus on its implicit claim to this interface role.

While the interface metaphor is undoubtedly somewhat misleading if taken literally as a claim about psycholinguistic processes (cf. the discussion of psycholinguistic implications of theories invoking "modularity" in Rapp and Goldrick 2000), it provides a useful way of thinking about the respective roles of *symbolic or discrete* representations and *parametric or continuous* ones in the description of language (e.g. Pierrehumbert 1990, Kornai 1994). I take it as uncontroversial that any detailed scientific description of physical processes must eventually be expressed in quantitative parametric terms. If that premise is accepted, then systematic phonetics can be interpreted as a hypothesis about the level of phonetic description *beyond which the use of symbolic representations ceases to be instructive or faithful to the empirical data*. In this light, Bloomfield's views quoted above become a competing hypothesis, namely that the level of description beyond which continuous parametric models are required is the *phonemic representation*, and that all other details of utterance phonetics cannot usefully be described in terms of symbolic categories. That is, regardless of whether the interface metaphor is ultimately enlightening psycholinguistically, there is an empirical issue here: is a symbolic idealization at the systematic phonetic level of description an appropriate part of an adequate scientific account of the sounds of language? A number of recent findings suggest that it is not.

The clearest evidence involves processes like assimilation, reduction and neutralization. In most conceptions of phonology, these are attributed to the workings of the phonological grammar – that is, they are part of *langue*. For example, vowels before coda nasals in English are routinely said to be allophonically nasalized: one symbolic abstraction (nasal vowel) is substituted for another (oral vowel). What we actually find in the instrumental record, though, is that the nasal airflow gradually increases across the vowel, quite unlike what happens in distinctively nasal vowels in a language like French (Cohn 1993). This means that any representation in which the vowel phone is categorically represented as either nasal or non-nasal fails to express

the difference between the phonetics of English and the phonetics of French. Conceivably the difference could be expressed in a systematic phonetic representation that allowed the transcriber to indicate different degrees of features like nasality, as suggested by Chomsky and Halle in *SPE* (1968: 65). However, that still precludes representing the time course of the velic opening, or any difference between the ranges of variability in the two languages. Similar comments apply to the usual conception of assimilations and deletions/reductions in connected speech, which are routinely represented as categorically either occurring or not occurring, and considered to be the output of the phonological grammar, e.g. /tɛn pæst tu/ → [tɛmpæstu]. A great many studies since the mid-1980s make it clear that such representations are an extremely crude reflection of the phonetic facts (e.g. Browman and Goldstein 1986, Nolan 1992, Zsiga 1997); there are many intermediate realizations, and it seems unlikely that sharp boundaries can be established between one categorical phone-based representation and another.[7]

These cases are directly relevant to the place of systematic phonetics within *langue*. In the view made explicit in *SPE,* the phonological grammar generates a detailed (but still symbolic and segmental) phonetic representation that contains complete information about assimilations and neutralizations and the like. This detailed phonetic representation is what is passed on to the physical realization system. The new phonetic findings suggest an alternative view: the grammar generates a rather more abstract interface representation – one that does not include any of the connected speech effects – and the interface representation is then passed on to a rather more elaborate physical realization system that specifies most aspects of pronunciation that are not the basis of categorical lexical distinctions. In such a conception of the sound system of a language, in effect, the phonology plays a smaller role in the description, while the role of phonetics is greater. The overall goal remains the same – to account for the fact that elements can count as the same in *langue* while exhibiting considerable systematic variability in their physical manifestations – but the interface between the symbolic system and the physical system is located in a different place. It seems reasonable to suggest that Bloomfield might have espoused such a view; more specifically, it seems that something like a classical phonemic transcription might serve as the "rather more abstract interface representation" that such a view requires.

None of the foregoing should be taken to suggest that the interface issue is purely a matter of efficient modelling or scientific description, devoid of psycholinguistic implications. On the contrary, the idea that the boundary between phonology and

[7] While this statement is certainly true of some connected speech processes, it is probably premature to conclude that all such processes involve gradiently variable output. A number of recent studies on a number of different languages suggest that it may be phonetically meaningful to distinguish sharply between assimilated and non-assimilated realizations in connected speech (e.g. Ellis and Hardcastle 2002, Ladd and Scobbie 2003, Kochetov and Pouplier 2008, Kainada 2009). It is not clear whether these apparently categorical effects in connected speech are related to purely phonetic "quantal" effects (Stevens 1972) or whether they reveal something important about the nature of the interface between the linguistic and the physical. Furthermore, the existence of gradiently variable connected speech effects does not preclude the possibility that such effects may become phonologized through language change. For example, Zsiga 1995 shows that the "assimilation" of /s/ to /ʃ/ before /j/ in English *confess your* is phonetically distinct from that in *confession,* suggesting that the /ʃ/ in *confession* is in some sense generated by the grammar whereas that in *confess your* is created by the workings of the physical realization system.

phonetics involves a representation less detailed than a systematic phonetic one is strengthened by evidence from what has been called covert contrast. First-language acquisition data is often said to involve neutralizations of adult contrasts, and various so-called phonological disorders involve children's alleged failure to distinguish adult phonemes, e.g. velar and coronal stops. However, detailed instrumental investigation (e.g. Macken and Barton 1980, Scobbie et al. 2000) suggests that in such cases children are sometimes – perhaps usually – aware of the phonological distinction and actually produce distinct patterns of articulation which, however, are not perceived as distinct by adults (including phone-based transcribers). The contrast is thus present in the child's *phonology*, but covertly, hidden from the observer equipped only with native speaker perceptual categories. In many cases it is impossible to characterize the way the child manifests the contrast in segment-based terms, but only in terms of continuous quantitative parameters.

The case of covert contrast shows that conceiving of systematic phonetics as an interface representation has concrete consequences for our understanding of developmental disorders and for the design of appropriate therapies. If the mapping from underlying representations to phones is part of *langue*, then children's phonological disorders are appropriately named, and physical phonetics – the motor behavior involved in realizing the linguistically specified output – is irrelevant theoretically and therapeutically. However, since it appears that phonologically disordered children are actually aware of the linguistic distinction and are unable to master the appropriate motor control to produce distinguishable acoustic output, then therapy obviously needs to focus on the physical, not the linguistic. And this, once again, suggests that the level of description corresponding to a systematic phonetic representation is not the right place to locate the interface between the categorical and the continuous.

3.3. What systematic phonetics could be a theory of

In order to avoid a potential misunderstanding, I should make clear that my remarks here are not intended as a blanket rejection of the IPA enterprise. As a tool for linguistic typology, systematic phonetics has an important role to play: terms like "front rounded vowel" and "uvular fricative" have reasonably clear language-independent definitions, and it is certainly meaningful to say that French and German have front rounded vowels while English and Spanish don't. Given what we now know about phonetic variability, statements like these must presumably be interpreted in something like the following way: French and German have phonological elements whose typical or canonical phonetic realization is a front rounded vowel, whereas English and Spanish do not, and any portion of an English or Spanish sentence that might be classed phonetically as a front rounded vowel is to be interpreted as the realization of some other phonological element. But whatever refinements of interpretation we wish to introduce into our understanding of phonetic typology, I believe that statements of this sort are useful scientific generalizations *about languages*. The problems discussed in the preceding two subsections arise from trying to use systematic phonetic terminology and concepts as descriptions of *individual acts of speech*.

The difficulty here is part of a more general problem with linguistic typology, better known from attempts to give language-independent definitions of parts of speech and

of grammatical notions like *subject*. The relation between definitions that are useful for typology and those that are needed for the description of individual words and constructions in individual languages has been discussed in a number of recent articles (e.g. Huddleston and Pullum 2002: 31f; Newmeyer 2007; Haspelmath 2007), and the way forward is still far from clear. I have elsewhere discussed the same issue in connection with the transcription of intonation (Ladd 2008a: sec. 3.2.2; 2008b), where the substantial current disagreements revolve in part around the nature of symbolic transcriptions. The point I wish to make here is simply that while systematic phonetics is of doubtful validity as the theoretical basis for describing utterance phonetics, it may be useful and important as a theory of phonetic typology. It is probably true, as noted by Pierrehumbert et al., (2000: 285), that "there are no two languages in which the implementation of analogous phonemes is exactly the same". It does not follow that systematic phonetic descriptions have no use anywhere in a scientific account of language; indeed, the very notion of "analogous phonemes" probably depends on such descriptions.

3.4. Segmental, suprasegmental, autosegmental

Although the discussion so far has focused almost exclusively on the phone, no discussion of phonetics in phonology would be complete without at least mentioning the problem of phonetic properties that fall outside the segmental idealization. Some such notion as "suprasegmental" or "prosodic" properties of speech has been assumed at least since the beginning of the IPA. However, its theoretical basis is clearly shaky. It takes little thought to realize that the traditional set of suprasegmental features – stress, pitch and quantity – are quite distinct phonetically. It is not much of an exaggeration to say that suprasegmentals are most accurately defined as those features that are not normally (or perhaps, not easily) represented in a segmental phonetic transcription. As such, they are effectively a by-product of the phone idealization.

The problem of defining suprasegmentals is discussed by Lehiste in the introduction to her influential book entitled simply *Suprasegmentals* (1970). She concedes that the conventional denotation of the term – stress, pitch and quantity – is essentially only a list, and notes that "a definition is preferable to a list". She then briefly identifies three main approaches to providing a real definition. The three are based on: (a) the fact that suprasegmental features can apply over domains longer than a segment; (b) the supposed phonetic distinctness of suprasegmental features from the properties that define segmental phones; (c) the supposed need to define suprasegmentals phonetically in terms of a syntagmatic comparison within an utterance (Jakobson, Fant and Halle 1952: 13). All of these have some element of usefulness and all fail in important ways; a full discussion of this topic must be left for a different paper.

In the 1970s serious attempts were made to understand suprasegmental phonology, stimulated by consideration of issues that were brought to light by the formalization of phonology in *SPE.* As noted above, the *SPE* formalism treated utterances as ordered strings of segments and segments as bundles of unordered features. The only place for stress, pitch and quantity in this formalization was as features of specific segments. Liberman (1975) and Liberman and Prince (1977) proposed that stress could more insightfully be treated in terms of a hierarchical "metrical" structure, which necessitates adding some sort of bracketing or constituent structure to Chomsky and Halle's simple strings. Specifically with regard to stress, this proposal

avoids some of the problems of phonetic interpretation that accompanied Chomsky and Halle's use of stress features on specific segments (see e.g. Vanderslice and Ladefoged 1972 for a typical reaction to the *SPE* analysis of stress). Much more fundamentally, the metrical proposal has led to a variety of theoretical ideas about constituent structure in phonology (e.g. Selkirk 1984, Nespor and Vogel 1986, Pierrehumbert and Beckman 1988) whose potential has, in my opinion, only begun to be explored. What seems clear, however, is that no unified set of suprasegmental features emerges from the metrical perspective; rather, it appears that stress is a very different kind of phenomenon from tone and quantity (cf. Ladd 2008, ch. 8).

About the same time that metrical phonology emerged as a response to the *SPE* treatment of stress, dissertations by Leben (1973) and Goldsmith (1976) tackled problems in the phonology of tone, leading to what came to be known by Goldsmith's term "autosegmental" phonology. Leben had demonstrated clearly that many ordinary phonological phenomena in tone languages are impossible to accommodate in any formalization of phonology that treats tone as a feature of a specific segment. Goldsmith pinpointed the problem as being what he called the "absolute slicing hypothesis", the idea that the signal can be exhaustively segmented into elements that succeed one another in time. He proposed instead that tones are an instance of a new kind of phonological element – an "autosegment" – that can be located on a separate "tier" from other segments, and that within the separate tier a separate set of temporal ordering relations obtains.[8]

In principle, the notion of autosegment could have a purely formal phonological definition. Specifically, what is noteworthy about tones from the point of view of the *SPE* formalism is that they are *not linearly ordered* with respect to segmental phonemes. This could be seen as a purely abstract mathematical property: *SPE*-style phonological strings are totally ordered, but phonological strings that allow for autosegments are only partially ordered (Ladd 2007). However, just as Trubetzkoy depended on concrete phonetic features to describe abstract phonological oppositions, so in developing the phonological abstraction of tiers Goldsmith focused on the fact that the *phonetic realization* of tone is not synchronized in lockstep with the phonetic realization of segments. Once this phonetic fact took center stage, it became obvious that it applies to almost any aspect of phonetic realization, and theoretical discussion within autosegmental phonology rapidly moved on to a consideration of the coordination of phonetic events in real time. Any special phonological properties of tone of the sort that concerned Leben were submerged beneath the idea that tone behaves like any other feature for purposes of synchronization, and the exploration of how feature tiers are temporally coordinated was extended to cover essentially phonetic phenomena such as assimilation as well.

The problem with this development is one we have already alluded to: autosegmental representations are discrete symbolic representations, and are poorly adapted to describing physical events. More generally, the extension of autosegmental phonology to deal with issues such as assimilation illustrates again the field's repeated failure to separate – really separate – phonetics and phonology: it appears that many phonologists want their descriptions to account for the phonetic detail of

[8] Precursors to the autosegment notion were discussed by Hockett 1955, especially sections 26 and 3222.

utterances. Yet most are reluctant to consider the use of formalisms involving continuous mathematics and quantitative variables, and without such formalisms, it is doubtful that any theory can deal adequately with all aspects of the linguistic use of sound.

4. Where do we go from here?

Early 21st century mainstream phonology, represented by Optimality Theory, has radically changed the form of its phonological grammar from the *SPE* codification, but continues to assume that the output of the grammar is a symbolic systematic phonetic representation, based on the phone concept. The broad continuity from the early days of the phonemic principle to the present is clear: despite the rather substantial theoretical upheavals of the early 1960s and the early 1990s, little has changed in the way most phonologists conceive of the interface between language and speech. However, as I have argued at some length, phone-based idealizations of speech are increasingly difficult to reconcile with the findings of phonetic research. We should not be surprised, then, that in the past couple of decades there have been a number of radical responses to the growing mismatch between phonological theorizing and empirical results in phonetics.

One response is to reject "formal phonology". This is the explicit proposal of a polemical article by Port and Leary 2005, who blame generative views about language as a discrete formal system for the idea that "phonetic segments are formal symbol tokens". The diagnosis here is faulty: treating phonetic segments as symbol tokens is, as we have seen, the essence of IPA transcription and of 20th century phonology generally. All that early generative phonology did was to formalize widely held views about phonetics. It may be appropriate to criticize formal phonology for many things, but it is not valid to treat it as the source of the phone concept.

Another more radical response is, in effect, to reject phonology altogether. This is the upshot of some versions of what is often known as "exemplar theory" (Goldinger 1996, Coleman 2002). The strong version of exemplar theory proposes that lexical entries are directly encoded in memory on the basis of acoustic traces, thereby bypassing the need for any representation in terms of phonological categories at all. This idea has a respectable pedigree (e.g. Klatt 1979) and seems likely to form part of an eventual fully worked-out psycholinguistic understanding of how words are represented in the mental lexicon (see further Coleman 2003). However, there is experimental evidence that makes clear that some such phonological abstraction as the phoneme is needed: perceivers can rapidly update their acoustic memory *of individual phonemes*, not only of whole words (McQueen et al. 2006). Updatable phoneme-sized categories form part of the modified exemplar theories espoused by e.g. Bybee 2001 and Pierrehumbert 2003; they accept the idea that fine phonetic detail is involved in lexical representations in some way, but they do not reject phonology altogether.

Within phonology, more or less the opposite response to exemplar theory is to argue for an outright divorce from phonetics. This case has been put most strongly in the recent literature by Hale and Reiss (e.g. 2000), who talk of "substance-free

phonology". The attractiveness of this view is that it takes Trubetzkoy's radical talk of the distinction between phonology and phonetics seriously, as Trubetzkoy himself did not. While a substance-free *phonology* may be possible and even desirable, though, one problem with this proposal is that it shows no interest in accounting for *language-specific phonetic* facts. That is, Hale and Reiss may be right (and Port and Leary wrong) that it is both possible and desirable to idealize language – *langue* – as a discrete formal system. However, a complete description of any actual language will always have to include statements about the language-specific interfaces between the formal system and the physical world. This is true both in semantics and in phonetics. In that sense no complete description can ever be "substance-free".

Another different approach to the growing mismatch between theories based on systematic phonetics and the results of research in laboratory phonology is that of so-called Articulatory Phonology (AP) (Browman and Goldstein 1986, 1989, Gafos 2002, and many others). AP represents a fundamental rethink of the interface notion and of the idea that phonological elements are symbolic abstractions: its elements are gestures, which are inherently quantitative abstractions and thus ideally suited to expressing the non-segmental aspects of phonetic realization that are increasingly coming to light. It seems quite clear that AP is correct in using a quantitative rather than a symbolic idealization *of phonetics*; time and physical space need to be modelled with continuous parameters, not categorical features or phones. Whether this also entails importing reference to actual time into our phonological abstractions, or getting rid of the segment-sized phoneme as an abstract element of the language, system is less clear; one could imagine marrying the AP approach to phonetic realization with some sort of "substance-free" phonology, in which case some sort of interface representation more abstract than a systematic phonetic representation would be required (cf. Ladd 2006). In any case, AP has so far shown little interest in accounting for some of the symmetries in phonological patterning that are so central to the tradition that begins with Trubetzkoy.

It is thus doubtful that any one of these new approaches by itself indicates the true path to enlightenment in the phonetic sciences, but collectively they all suggest that a fundamental re-evaluation of the place of systematic phonetics in phonology is overdue. We have spent too much time as a field failing to resolve foundational issues and worrying about problems that simply disappear if seen in a different light. Unexamined acceptance of the twin assumptions of systematic phonetics – the segmental idealization and the universal categorization assumption – is certainly partly to blame, and recognizing systematic phonetic theory as an important source of confusion in phonology provides us with an opportunity to make real progress.

A number of things begin to make sense if we assume that there is no closed universal set of phonetic elements out of which utterances are built. For example, during the 1970s and 1980s it was widely supposed that infants are born with the ability to perceive all possible phonemic contrasts and gradually lose the ones they don't need in their language. We now know that adults preserve the ability to perceive some differences that are non-phonemic in their language (e.g. Best, McRoberts and Goodell 2001), that children make various false steps (e.g. Pater, Stager and Werker 2004), and more generally that children have to *learn* certain distinctions of their native language. This is exactly what we should expect if there is no universally valid

categorization of phonetic segments, because without that categorization, the very concept "all possible contrasts" is incoherent.

Another theoretical conundrum that becomes suspect once we start questioning the validity of systematic phonetics is the problem of opacity in so-called chain shifts in first and second language acquisition, the famous "puddle puggle puzzle" (Smith 1973). There are many well-known cases in which language acquirers are said to replace [A] by [B] but at the same time replace [C] by [A], such as Smith's original example of /d/ → /g/ and /z/ → /d/, or the somewhat more complicated case of /s/ → [θ], /θ/ → [f], /f/ → [f] cited by Dinnsen and Barlow 1998. But these cases are problematical only insofar as the identification of phones is accurate. The second case is a problem only if (a) [f] < /θ/ is identical to [f] < /f/ in the child's speech (which is doubtful given the existence of covert contrast) and (b) [θ] < /s/ in the child's speech can reliably be identified with [θ] < /θ/ in adult speech (which is uncertain at best). Similarly, Smith's classic case is a problem only if [d] < /z/ in the child's speech can be identified with [d] < /d/ in adult speech. If the phonetic realizations are actually empirically distinct, the chain shift problem evaporates.

But much more central issues are at stake. The most conspicuously unresolved issue in phonology, in my view, is the debate over the classical phoneme that began in the late 1950s. The critiques by Halle 1959 and Chomsky 1964 deprived the traditional phoneme concept of its theoretical legitimacy, but it has nevertheless survived more or less intact for the intervening half-century, in practical applications (such as speech therapy, reading and literacy training, and speech technology), in linguistic fieldwork, and – revealingly – in beginning linguistics courses.[9] Schane's contention (1971) that generative phonology had only superficially done away with the phoneme has never been refuted (see further Ladd 2006).

Within the generative tradition, the problem of defining the classical phoneme manifests itself as the "abstractness controversy", first raised in such terms by Kiparsky 1968. By comparison to classical phonemic analyses, *SPE*-style systematic phonemic representations tend to be "more abstract", in the specific sense that they differ more from the corresponding string of phones. Kiparsky pointed to various undesirable consequences of allowing unlimited abstractness in this sense, but given the *SPE* framework he was unable to find many principled reasons for avoiding it. The abstractness problem is inherent in the *SPE* architecture: "systematic phonemes" and phones are the same kind of formal object, namely bundles of features, and it is difficult to constrain a set of ordered feature-changing rules except *ad hoc*. From the *SPE* point of view, that is, classical phonemic theory amounted to little more than a collection of arbitrary restrictions on permissible phonological abstractions – yet there seemed to be no non-arbitrary basis for a different set of restrictions.

The controversy based on Kiparsky's statement of the problem was an active topic of debate for several years (e.g. Jackendoff 1975, Lightner 1975) but was never resolved. Instead, in the late 1970s it was merely put aside, as bright young theorists started working on other problems, in particular autosegmental and metrical phonology. Lexical Phonology in the 1980s (e.g. Halle and Mohanan 1985, Kaisse

[9] A web search on the set of terms {*allophone, phoneme, phonology, introduction*} in March 2008 yielded more than 75,000 hits; many of the first hundred hits are handouts or syllabuses from beginning linguistics courses from major universities, including one from MIT's open courseware.

and Shaw 1985) was an attempt to deal with some of the problems Kiparsky had discussed, but in a certain sense did no more than rehabilitate something like a classical phonemic representation without resolving the question of the phoneme's theoretical legitimacy, and has accordingly failed to live up to its apparent promise. Within OT, the abstractness issue has not been explicitly discussed, presumably because of the notion that OT does not involve derivations or rules. However, the problem is still present, because the OT formalism is like the SPE formalism in the sense that it provides a way of establishing correspondences between one symbolic representation and another. The fact that OT has not addressed the question of what the abstract "input" representations are for or how they should be determined does not mean that this is not a problem; the principal acknowledgement of this issue has been in the form of proposals to consider "output-output" constraints (e.g. Benua 2000, Burzio 2002). An OT based purely on output-output constraints, with due attention paid to the nature of the output representation, would represent a genuinely radical departure from SPE, but so far that has not happened.

The theoretical issues surrounding the phoneme and the abstractness of phonological representations have always been framed in the way they are because of the assumption that the point of the exercise is to map one symbolic abstraction onto another: phonemes onto phones, systematic phonemic onto systematic phonetic representation, OT input onto OT output. My goal in this paper has been to show that this assumption is a direct consequence of having based phonological theory on the IPA theory of systematic phonetics. If instead we start from the assumption that phonetic realization involves a mapping from symbolic phonological abstractions of some sort to a *continuous* signal describable in *quantitative* physical terms, then one of our primary theoretical tasks as phonologists must be to clarify the nature of those abstractions – in effect, to define the phoneme. Systematic phonetics almost certainly has a useful scientific role to play in an overall understanding of language. But I don't believe that we will make much progress in phonology until we stop trying to ground our theories in the systematic phonetic representation of individual acts of speech.

Acknowledgements

Work on this paper was supported in part by an Individual Research Fellowship ("Simultaneous and Sequential Structure in Language") from the Leverhulme Trust. I thank Wayles Browne, Patrick Honeybone, John Joseph, Mits Ota, Geoff Pullum, Barbara Scholz, and Donca Steriade for comment, conversation, and pointers to the literature.

References

Archangeli, Diana, and D. Terence Langendoen (1997). *Optimality Theory: An Overview.* Oxford: Blackwell.

Baltaxe, Christiane (translator) (1969). *Principles of Phonology.* (English translation of Trubetzkoy 1958). Berkeley and Los Angeles: University of California Press.

Beckman, Mary E. and John Kingston (1990). Introduction. In Kingston and Beckman, pp. 1-16.

Benua, Laura (2000). *Phonological Relations between Words.* Abingdon, Oxon: Routledge.

Best, Catherine T., Gerald W. McRoberts and Elizabeth Goodell (2001). Discrimination of non-native consonant contrasts varying in perceptual assimilation to the listener's native phonological system. *Journal of the Acoustical Society of America* 109: 775-794

Bloch, Bernard (1941). Phonemic overlapping. *American Speech* 16: 278-284. Reprinted in Joos 1966, pp. 93-96.

Bloch, Bernard (1948). A set of postulates for phonemic analysis. *Language* 24: 3-46. Reprinted in Makkai 1972, pp. 167-199.

Bloomfield, Leonard (1933). *Language.* New York: Holt, Rinehart & Winston.

Boersma, Paul (1998). *Functional Phonology. Formalizing the interactions between articulatory and perceptual drives.* Doctoral dissertation, University of Amsterdam.
The Hague: Holland Academic Graphics.

Browman, Catherine P. and Louis Goldstein (1986). Towards an articulatory phonology. *Phonology Yearbook* 3: 219-252.

Browman, Catherine P. and Louis Goldstein (1989). Articulatory gestures as phonological units. *Phonology* 6: 201-251.

Burzio, Luigi (2002). Segmental contrast meets output-to-output faithfulness. In G. Fanselow and C. Féry (eds.), *Resolving Conflicts in Grammars: Optimality Theory in Syntax, Morphology, and Phonology* (Hamburg: Buske Verlag), 163-182.

Bybee, Joan (2001). *Phonology and Language Use.* Cambridge: Cambridge University Press.

Cantineau, Jean (translator) (1949). *Principes de Phonologie* (French translation of Trubetzkoy 1958). Paris: Klincksieck.

Chang Nien-Chuang (1958). Tones and intonation in the Chengtu dialect (Szechuan, China). *Phonetica* 2: 59–84.

Cho, Taehong and Peter Ladefoged (1999). Variation and universals in VOT: evidence from 18 languages. *Journal of Phonetics* 27: 207-229.

Chomsky, Noam (1964). The nature of structural descriptions. Chapter 4 of *Current Issues in Linguistic Theory* (The Hague: Mouton). Reprinted in Makkai 1972, pp. 401-423.

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

Cohn, Abigail (1993). Nasalisation in English: phonology or phonetics. *Phonology*10: 43-81.

Coleman, John (2002). Phonetic representations in the mental lexicon. In J. Durand, and B. Laks (eds.), *Phonetics, phonology and cognition* (Oxford: Oxford University Press), pp. 96–130.

Coleman, John (2003). Discovering the acoustic correlates of phonological contrasts. *Journal of Phonetics* 31: 351-372.

Dinnsen, Daniel A. and Jessica A. Barlow (1998). On the characterization of a chain shift in normal and delayed phonological acquisition. *Journal of Child Language* 25: 61-94.

Ebert, Karen (1975-1979). *Sprache und Tradition der Kera (Tschad)* (three volumes). Berlin: Reimer

Ellis, Lucy and William J. Hardcastle (2002). Categorical and gradient properties of assimilation in alveolar to velar sequences: evidence from EPG and EMA data. *Journal of Phonetics* 30: 373-396.

Flemming, Edward and Stephanie Johnson (2007). Rosa's roses: Reduced vowels in American English. *Journal of the International Phonetic Association* 37: 83-96.

Gafos, Adamantios (2002). A grammar of gestural coordination. *Natural Language and Linguistic Theory* 20: 269-337.

Goldinger, S. D. (1996). Words and voices: Episodic traces in spoken word identification and recognition memory. *Journal of Experimental Psychology: Learning, Memory and Cognition* 22: 1166–1183.

Goldsmith, John (1976). *Autosegmental Phonology*. PhD dissertation, MIT. Published 1979 by Garland Press, New York.

Hale, Mark and Charles Reiss (2000). Substance abuse and dysfunctionalism: Current trends in phonology. *Linguistic Inquiry* 31: 157-169.

Halle, Morris (1959). *The Sound Pattern of Russian.* The Hague: Mouton.

Halle, Morris, and K. P. Mohanan (1985). Segmental phonology of Modern English. *Linguistic Inquiry* 16: 57-116.

Haspelmath, Martin (2007). Pre-established categories don't exist: Consequences for language description and typology. *Linguistic Typology* 11: 119–32.

Hayes, Bruce and Donca Steriade (2004). Introduction: the phonetic bases of phonological Markedness. In B. Hayes, R. M. Kirchner, and D. Steriade (eds.), *Phonetically Based Phonology* (Cambridge: Cambridge University Press), pp. 1-33.

Hockett, Charles F. (1942). A system of descriptive phonology. *Language*18: 3-21. Reprinted in Joos 1966, pp. 97-108.

Hockett, Charles F. (1955). *A Manual of Phonology.* Baltimore: Waverley Press.

Huddleston, Rodney and Geoffrey K. Pullum (2002). *The Cambridge Grammar of the English Language.* Cambridge: Cambridge University Press.

Hyman, Larry M. (1975). *Phonology: Theory and Analysis.* New York: Holt, Rinehart & Winston.

IPA (1949). *The Principles of the International Phonetic Association.* London: University College.

IPA (1999). *Handbook of the International Phonetic Association.* Cambridge: Cambridge University Press.

Jackendoff, Ray S. (1975). Morphological and Semantic Regularities in the Lexicon, *Language* 51: 639-671.

Jakobson, Roman and Morris Halle (1956). *Fundamentals of Language.* The Hague: Mouton.

Jakobson, Roman, Gunnar Fant, and Morris Halle (1952). *Preliminaries to Speech Analysis.* The Hague: Mouton.

Joos, Martin (1948). *Acoustic Phonetics* (*Language Monograph* 23). Baltimore: Waverley Press.

Joos, Martin (1966). *Readings in Linguistics I.* Chicago: University of Chicago Press.

Kager, René (1999). *Optimality Theory.* Cambridge: Cambridge University Press.

Kainada, Evia (2009). *Effects of prosodic boundary strength on segmental and suprasegmental processes in Modern Greek.* PhD dissertation, University of Edinburgh.

Kaisse, Ellen and Patricia Shaw (1985). On the theory of Lexical Phonology. *Phonology Yearbook* 2: 1-30.

Kenstowicz, Michael and Charles Kisseberth (1979). *Generative Phonology.* Orlando: Academic Press.

Kingston,, John and Mary E. Beckman (eds.) (1990). *Papers in Laboratory Phonology I: Between the Grammar and Physics of Speech.* Cambridge: Cambridge University Press.

Kiparsky, Paul (1968). How abstract is phonology? Paper distributed by Indiana University Linguistics Club.

Klatt, Dennis (1979). Speech perception: A model of acoustic-phonetic analysis and lexical access. *Journal of Phonetics* 7: 279-312.

Kochetov, Andrei and Marianne Pouplier (2008). Phonetic variability and grammatical knowledge: an articulatory study of Korean place assimilation. *Phonology* 25:

Kornai, András (1994). Relating phonetic and phonological categories. In E. S. Ristad (ed.), *Language Computations,* vol. 17 (Philadelphia: American Mathematical Society), pp. 21-36.

Ladd, D. Robert (2006). "Distinctive phones" in surface representation. In L. Goldstein, D. H. Whalen, C. T. Best (eds.), *Labgoratory Phonology 8* (Berlin: Mouton de Gruyter), pp. 1-26.

Ladd, D. Robert (2007). Tone, autosegmental phonology, and the partial ordering of phonological elements. Paper presented at the CUNY Forum on Precedence Relations in Phonology, January 2007.

Ladd, D. Robert (2008a). *Intonational Phonology* (second edition). Cambridge: Cambridge University Press.

Ladd, D. Robert (2008b). Review of Sun-Ah Jun (ed.), *Prosodic Typology* (Oxford University Press, 2005). *Phonology* 25: 372-376.

Ladd, D. Robert and James M. Scobbie (2003). External sandhi as gestural overlap? Counter-evidence from Sardinian. In J. Local, R. Ogden and R. Temple (eds), *Papers in Laboratory Phonology VI: Phonetic Interpretation* (Cambridge: Cambridge University Press), pp. 164-182.

Ladefoged, Peter (1990). Some reflections on the IPA. *Journal of Phonetics* 18: 335-346.

Laver, John (1994). *Principles of phonetics.* Cambridge: Cambridge University Press.

Leben, William (1973). *Suprasegmental Phonology*. PhD dissertation, MIT. Published 1980 by Garland Press, New York.

Lehiste, Ilse (1970). *Suprasegmentals.* Cambridge MA: MIT Press.

Liberman, Mark Y. (1975). *The Intonational System of English.* PhD dissertation, MIT. Published 1979 by Garland Press, New York.

Liberman, Mark Y. and Alan Prince (1977). On stress and linguistic rhythm. *Linguistic Inquiry* 8: 249–336.

Lightner, Theodore M. (1975). The role of derivational morphology in generative grammar. *Language* 51: 617-638.

Macken, Marlys and David Barton (1980). A longitudinal study of the acquisition of the voicing contrast in American-English word initial stops, as measured by voice onset time. *Journal of Child Language* 7: 41-72.

Makkai, Valerie Becker, ed. (1972) *Phonological Theory: Evolution and Current Practice.* New York: Holt, Rinehart & Winston.

Martinet, André (1939). Un ou deux phonèmes? *Acta Linguistica* 1: 94-103. Reprinted in E. Hamp, F. W. Householder, and R. Austerlitz (eds.), *Readings in Linguistics II* (Chicago: University of Chicago Press, 1966), pp. 116-123.

McQueen, James, Anne Cutler, and Dennis Norris (2006). Phonological abstraction in the mental lexicon. *Cognitive Science* 30: 1113-1126.

Nearey, Terrance M. (1990). The segment as a unit of speech perception. *Journal of Phonetics* 18: 347-373.

Nespor, Marina and Irene Vogel (1986). *Prosodic Phonology.* Dordrecht: Foris.

Newmeyer, Frederick J. (2007). Linguistic typology requires crosslinguistic formal categories. *Linguistic Typology* 11: 133–57.

Nolan, Francis (1992). The descriptive role of segments: evidence from assimilation. In G. J. Docherty and D. R. Ladd (eds.), *Papers in Laboratory Phonology II: Gesture, Segment, Prosody.* Cambridge: Cambridge University Press.

Odden, David (1994). Adjacent parameters in phonology. *Language* 70: 289-330.

Pater, Joe, Christine Stager, and Janet Werker (2004). The perceptual acquisition of phonological contrasts. *Language* 80: 384-402.

Pearce, Mary (2007). *The interaction of tone with voicing and foot structure: Evidence from Kera phonetics and phonology.* PhD dissertation, University College London.

Pierrehumbert, Janet B. (1990). Phonological and phonetic representation. *Journal of Phonetics* 18: 375-394.

Pierrehumbert, Janet B. (2003) Probabilistic phonology: Discrimination and robustness. In R. Bod, J. Hay, and S. Jannedy (eds.) *Probabilistic Linguistics* (Cambridge MA: MIT Press), pp. 177-228.

Pierrehumbert, Janet and Mary E. Beckman (1988). *Japanese Tone Structure*. Cambridge, MA: MIT Press.

Pierrehumbert, Janet B., Mary E. Beckman, and D. Robert Ladd (2000). Conceptual foundations of phonology as a laboratory science. In N. Burton-Roberts, P. Carr, and G. J. Docherty (eds.), *Phonological Knowledge: Conceptual and Empirical Issues* (Oxford: Oxford University Press), pp. 273-303.

Pike, Kenneth (1943). *Phonetics.* Ann Arbor: Univ. of Michigan Press.

Port, Robert F. and Adam P. Leary (2005). Against formal phonology. *Language* 81: 927-964.

Postal, Paul M. (1968). *Aspects of Phonological Theory.* New York: Harper & Row.

Prince, Alan and Paul Smolensky (2004). *Optimality Theory in Phonology: A Reader.* Oxford: Blackwell.

Rapp, Brenda and Matthew Goldrick (2000). Discreteness and Interactivity in Spoken Word Production. *Psychological Review* 107: 460-499.

Rose, Sharon and Rachel Walker (2004). A typology of consonant agreement as correspondence. *Language* 80: 475-531.

Schane, Sanford A. (1971). The phoneme revisited. *Language* 47: 503-521.

Schane, Sanford A. (1973). *Generative Phonology.* Englewood Cliffs NJ: Prentice-Hall.

Scobbie, James M., Fiona Gibbon, William J. Hardcastle, and Paul Fletcher (2000). Covert contrast as a stage in the acquisition of phonetics and phonology. In M. Broe and J. Pierrehumbert (eds.) *Papers in Laboratory Phonology V: Language Acquisition and the Lexicon* (Cambridge: Cambridge University Press), pp. 194-207.

Selkirk, Elisabeth O. (1984). *Phonology and syntax: the relation between sound and structure.* Cambridge, MA: MIT Press.

Smith, Neil (1973). *The acquisition of phonology*. Cambridge: Cambridge University Press.

Stevens, Kenneth N. (1989). On the quantal nature of speech. *Journal of Phonetics* 17: 3-45.

Trubetzkoy, Nikolai Sergeevich (1958). *Grundzüge der Phonologie.* Göttingen: Vandenhoeck & Ruprecht.

Vanderslice, Ralph and Peter Ladefoged (1972). Binary suprasegmental features and transformational word-accentuation rules. *Language* 48: 819–38.

Zsiga, Elizabeth C. (1995). An acoustic and electropalatographic study of lexical and post-lexical palatalization in American English. In B. Connell and A. Arvaniti (eds.), *Papers in laboratory phonology IV: Phonology and phonetic evidence* (Cambridge: Cambridge University Press), pp. 282-302.

Zsiga, Elizabeth C. (1997). Features, gestures, and Igbo vowels: An approach to the phonology-phonetics interface. *Language* 73: 227-274.

Zwirner, Eberhard and Kurt Zwirner (1966). *Grundfragen der Phonometrie* (second edition). Basel: S. Karger.