

In G. Fanselow, C Féry, R Vogel & M. Schlesewsky  
(2006) *Gradience in Grammar: Generative  
Perspectives*. Oxford: OUP pp. 25-44.

2

---

## Is there Gradient Phonology?

ABIGAIL C. COHN

### 2.1 Introduction

In this chapter,<sup>1</sup> I consider the status of gradient phonology, that is, phonological patterns best characterized in terms of continuous variables. I explore some possible ways in which gradience might exist in the phonology, considering the various aspects of phonology: contrast, phonotactics, morphophonemics, and allophony. A fuller understanding of the status of gradience in the phonology has broader implications for our understanding of the nature of the linguistic grammar in the domain of sound patterns and their physical realizations. In the introduction, I consider why there might be gradience in the phonology (Section 2.1.1). I then briefly discuss the nature of phonology versus phonetics (Section 2.1.2).

#### 2.1.1 *Is there gradient phonology?*

Phonology is most basically a system of contrasts, crucial to the conveyance of linguistic meaning. This suggests that phonology is in some sense 'categorical'. Central to most formal models of phonology is a characterization of minimally contrasting sound 'units' (whether in featural, segmental, or gestural terms) that form the building blocks of meaningful linguistic units. In what ways is phonology categorical—mirroring its function as defining contrast, and to what degree is phonology inherently gradient in its representation, production, perception, acquisition, social realization, and change over time?

<sup>1</sup> A number of the ideas discussed in this chapter were developed in discussions in my graduate seminars at Cornell, Spring 2004 and Spring 2005. Some of these ideas were also presented in colloquia at the Universities of Buffalo and Cornell. Thanks to all of the participants in these fora for their insightful comments and questions. Special thanks to Mary Beckman, Jim Scobbie, and an anonymous reviewer for very helpful reviews of an earlier draft, as well as Johanna Brugman, Marc Brunelle, Ioana Chitoran, Nick Clements, Caroline Féry, Lisa Lavoie, Amanda Miller, and Draga Zec for their comments.

- The physical realization of sounds, understood (at least intuitively) as abstract units, is continuous in time and space, with the relationship between the specific acoustic cues and abstract contrasts often being difficult to identify.
- One crucial aspect of the acquisition of a sound system is understanding how phonetic differences are marshalled into defining abstract categories.
- Intraspeaker and interspeaker variation signal speaker identity, community identity, and attitude, while simultaneously conveying linguistic meaning through minimally contrasting elements.
- The results of many diachronic changes, understood to be 'regular sound change' in the Neogrammarian sense, are categorical, yet how do changes come about? Are the changes themselves categorical and abrupt or do the changes in progress exhibit gradience and gradual lexical diffusion?

A modular view of grammar such as that espoused by Chomsky and Halle (1968, SPE) frames our modelling of more categorical and more gradient aspects of such phenomena as belonging to distinct modules (e.g. phonology versus phonetics). While SPE-style models of sound systems have achieved tremendous results in the description and understanding of human language, strict modularity imposes divisions, since each and every pattern is defined as either X or Y (e.g. phonological or phonetic). Yet along any dimension that might have quite distinct endpoints, there is a grey area. For example, what is the status of vowel length before voiced sounds in English, *bead* [bi:d] versus *beat* [bit]? The difference is greater than that observed in many other languages (Keating 1985), but does it count as phonological?

Bearing in mind how a modular approach leads to a particular interpretation of the issues, I consider the relationship between phonology and phonetics before exploring the question of gradience in phonology.

### 2.1.2 The nature of phonetics versus phonology

A widely held hypothesis is that phonology is the domain of abstract patterns understood to be discrete and categorical, and phonetics is the domain of the quantitative realization of those patterns in time and space. These relationships are sketched out in (2.1).

- (2.1) The relationship between phonology and phonetics:
- |           |   |                       |
|-----------|---|-----------------------|
| phonology | = | discrete, categorical |
| ≠         |   |                       |
| phonetics | = | continuous, gradient  |

For recent discussions of this consensus view, see for example Keating (1996); Cohn (1998); Ladd (2003), also individual contributions in Burton-Roberts *et al.* (2000) and Hume and Johnson (2001). See also Cohn (2003) for a fuller discussion of the nature of phonology and phonetics and their relationship.

For the sake of concreteness, consider an example of phonological patterns and their corresponding phonetic realization that are consistent with the correlations in (2.1). In Figure 2.1, we see representative examples of the patterns of nasal airflow in French and English (as discussed in Cohn 1990, 1993). Nasal airflow is taken here as the realization of the feature Nasal.

In the case of a nasal vowel in French, here exemplified in the form *daim* 'deer' [dɛ̃] (Figure 2.1a), there is almost no nasal airflow on [d] and there is significant airflow throughout the [ɛ̃]. Here we observe *plateaus* corresponding to the phonological specifications, connected by a rapid transition. In English on the other hand, during a vowel preceding a nasal consonant, such as [ɛ] in *den* [dɛn] (Figure 2.1b), there is a gradient pattern—or a *cline*—following the oral [d] and preceding the nasal [n] (which are characterized by the absence and presence of nasal airflow respectively). This is quite different from the pattern of nasalization observed on the vowel in cases like *sent* [sɛ̃t] (Figure 2.1c), in which case the vowel is argued to be phonologically nasalized (due to the deletion of the following /n/) and we observe a plateau of nasal airflow during the vowel, similar to the pattern seen in French. The observed




FIGURE 2.1. Examples of nasal airflow in French and English following Cohn (1990, 1993)

differences between French and English relate quite directly to the fact that French has nasal vowels, but English does not.

If the correlations in (2.1) are correct, we expect to find categorical phonology, but not gradient phonology, and gradient, but not categorical, phonetics. Recent work calls into question this conclusion. In particular, it is evidence suggesting that there is gradience in phonology that has led some to question whether phonetics and phonology are distinct. Pierrehumbert *et al.* (2000) state the question in the following way:

this assertion [that the relationship of quantitative to qualitative knowledge is modular] is problematic because it forces us to draw the line somewhere between the two modules. Unfortunately there is no place that the line can be cogently drawn. ... In short, knowledge of sound structure appears to be spread along a continuum. Fine-grained knowledge of continuous variation tends to lie at the phonetic end. Knowledge of lexical contrasts and alternations tend to be more granular. (Pierrehumbert *et al.* 2000: 287)

Let us consider the background of this issue in a bit more depth. Growing out of Pierrehumbert's (1980) study of English intonation, gradient phonetic patterns are understood as resulting from phonetic implementation, through a mapping of categorical elements to continuous events. Under the particular view developed there, termed *generative phonetics*, these gradient patterns are the result of interpolation through phonologically unspecified domains. Keating (1988) and Cohn (1990) extend this approach to the segmental domain, arguing that phenomena such as long distance pharyngealization and nasalization can be understood in these terms as well. For example, the cline in nasal airflow seen in the vowel [ɛ] in [dɛn] in Figure 2.1b is interpreted as resulting from phonetic interpolation through a phonologically unspecified span.

The phonology, then, is understood as the domain of discrete, qualitative patterns and the phonetics as the domain of the continuous, quantitative realization of those patterns. Intrinsic to this view is the idea that lexical entries and phonological patterns are represented in terms of distinctive features, taken to be abstract properties, albeit defined phonetically. These are then interpreted in a phonetic component, distinct from the phonological one. I refer to this as a *mapping approach*. A modular mapping approach has been the dominant paradigm to the phonology-phonetics interface since the 1980s and has greatly advanced our understanding of phonological patterns and their realization. Such results are seen most concretely in the success of many speech-synthesis-by-rule systems both in their modelling of segmental and suprasegmental properties of sound systems. (See Klatt 1987 for a review.)

An alternative to the types of approaches that assume that phonology and phonetics are distinct and that there is a mapping between these two modules or domains are approaches that assume that phonology and phonetics are one and the same thing, understood and modelled with the same formal mechanisms, what I term a *unidimensional* approach. A seminal approach in this regard is the theory of Articulatory Phonology, developed by Browman and Goldstein (1992 and work cited therein), where it is argued that the domains that are often understood as phonology and phonetics respectively can both be modelled with the same formalisms as constellations of gestures. Under this view, phonetics and phonology are not distinct and the apparent differences might arise through certain (never explicitly specified) constraints on the phonology. This gestural approach has served as fertile ground for advancing our understanding of phonology as resulting at least in part from gestural coordination. However, there are criticisms of this approach as a comprehensive theory of phonology, including arguments that Articulatory Phonology greatly overgenerates possible patterns of contrast. (See commentaries by Clements 1992 and Steriade 1990.)

More recently, there is a significant group of researchers (e.g. Flemming 2001; Kirchner 2001; Steriade 2001; see also Hayes *et al.* 2004) working within constraint-based frameworks, pursuing the view that there is not a distinction between constraints that manipulate phonological categories and those that determine fine details of the representation. This then is another type of unidimensional approach that assumes no formally distinct representations or mechanisms for phonology and phonetics. One type of argument in favour of this approach is that it offers a direct account of naturalness in phonology. However, the strength of this argument depends on one's view about the source(s) of naturalness in language. (See Blevins 2004 for extensive discussion of this issue.)

Such unidimensional views of phonology and phonetics also need to offer an account of not only what I term 'phonetics in phonology', but also of the relationship between phonological units and physical realities—'phonology in phonetics'. (See Cohn 2003 for a discussion of the distinct ways that phonology and phonetics interact with each other.) Independent of the account of naturalness, the question of whether one can adequately model the way that the phonetics acts on phonology still remains. Both Zsiga (2000) and Cohn (1998) have argued that such unidimensional approaches do not offer an adequate account. As documented by Cohn (1998), this is commonly seen in 'phonetic doublets', cases where similar but distinct effects of both a categorical and gradient nature are observed in the same language. These sorts of effects can be seen in the case of nasalization discussed above. In French, in

the realization of contrastive nasal vowels, there is nasal airflow resulting from the contrast and also from coarticulatory patterns, seen, for example, in the transition between oral vowels and nasal consonants. Both aspects need to be modelled. In the case of contextual nasalization in English, there are both long distance and more local effects seen in the physical patterns of nasal airflow that need to be accounted for.

The question of whether phonology and phonetics should be understood as distinct modules needs to be approached as an empirical question. What sort of approach gives us the best fit for the range of more categorical versus more gradient phenomena?

There are clearly some grey areas—notably gradient phonology. Yet it is important to realize that just because it is difficult to know exactly where to draw the line (cf. Pierrehumbert *et al.* 2000), this does not mean there are not two separate domains of sound structure. The fact that it is difficult to draw a line follows in part from the conception of *phonologization* (Hyman 1976), whereby over time low-level phonetic details are enhanced to become phonological patterns. Phonologization by its very nature may result in indeterminate cases. As phonetic details are being enhanced, it will be difficult at certain stages to say that a particular pattern is ‘phonetic’ while another is ‘phonological’. It has been suggested, for example that vowel lengthening before voiced sounds in English is currently in this in-between state. The difficulty of drawing a line also relates to the sense in which categoriality can only be understood in both rather abstract and language-specific terms.

Recent work suggests that phonology and phonetics are not the same thing, but that the distinction might be more porous than assumed following strict modularity (e.g. Pierrehumbert 2002 and Scobbie 2004). Pierrehumbert (2002: 103) states: ‘categorical aspects of phonological competence are embedded in less categorical aspects, rather than modularized in a conventional fashion.’ We return below to the nature of the relationship between phonology and phonetics, as the status of gradient phonology plays a crucial role in this question.

In order to investigate gradience in phonology, we need a clearer understanding of what we mean by gradience and we need to consider how it might be manifested in different aspects of the phonology. I turn to these questions in the next section.

## 2.2 Aspects of gradience

Most basically, we understand *gradient* and *gradience* in opposition to *categorical* and *categoriality*. A *gradient* (n.) in its original sense is a mapping

from one continuous variable to another, that is, a slope. (In linguistic usage, we use the form *gradience* as a noun and *gradient* as an adjective.) It has also shifted to mean the continuous nature of a single variable.<sup>2</sup> Thus we need to be clear on which sense of *gradient* we are talking about. Discrete is often equated with categorical and continuous with gradient (although there may be gradient patterns that are discrete). We need to consider both the question of what is gradient, as well as what is continuous.

The terms *gradient* and *gradience* have been used in a number of different ways in the recent phonetic and phonological literature. To think more systematically about the nature of gradience in phonology, we need to tease apart these different usages (Section 2.2.1) before considering how these senses might apply to different aspects of what is understood to be phonology—that is, contrast (Section 2.2.2), phonotactics (Section 2.2.3), and alternations, both morphophonemics (Section 2.2.4) and allophony (Section 2.2.5).

### 2.2.1 *Different uses of the term gradience*

When we talk about sound patterns, there are at least three senses of *gradience* that have been prevalent in the recent literature—temporal/spatial gradience, variability, and gradient well-formedness.<sup>3</sup>

**2.2.1.1 *Temporal/spatial gradience*** In work on the phonetic implementation of phonology, *gradient/gradience* is used in the sense of change in phonetic space through time. This is the sense of gradient versus categorical seen in the example of the realization of nasalization shown in Figure 2.1. In this case, there is a change in the amount of nasal airflow (space) through time, characterized as being a cline, distinct from more plateau-like cases (argued to obtain in cases of contrast). This is what I take to be the primary sense of gradience versus categoriality as it applies to the domain of sound patterns and their realization.

**2.2.1.2 *Variability*** The term *gradience* is also often used to refer to variable realizations or outcomes of sound patterns, understood as unpredictable or as stemming from various sociolinguistic and performance factors. We might understand this as gradience in the sense of gradience across tokens.

<sup>2</sup> Thanks to Mary Beckman (p.c.) for clarifying this question of usage.

<sup>3</sup> There is an additional use of the term *gradient* in the recent phonological literature. Within Optimality Theory, *gradient* has also been used to refer to constraint satisfaction (e.g. McCarthy and Prince 1993; McCarthy 2003), where more violations of a particular constraint are worse than a single violation. This is different from the other senses discussed here and will not come into play in the present discussion.

Variability is sometimes understood in phonological terms as optional rule application, or freely ranked constraints, or as 'co-phonologies'. These patterns have sometimes been modelled in statistical or stochastic terms. There are also approaches that model these factors directly as sociolinguistic or stylistic markers. (See Antilla 2002 and Coetzee 2004 for discussion of recent approaches to modelling phonological variation.) Both variability and gradience in phonetic implementation are pervasive in phonetic patterns and both must ultimately be understood for a full understanding of phonology and its realization.

What we sometimes interpret as variability may in fact result from methodological approaches that are not fine-tuned enough in their characterization of conditioning factors or prosodic context. For example, the realization of the contrast between so-called 'voiced' and 'voiceless' stops in English is highly dependent on segmental context, position in the word, position in the utterance, location relative to stress, etc. The nature of contrast may also vary systematically by speaker (Scobbie 2004). If these factors are not taken into consideration, one would conclude that there is enormous variability in the realization of these contrasts, while in fact much of the variation is systematic.

It is not necessarily the case that temporal/spatial gradience and variability go hand in hand. In fact, there are well documented cases where they do not, that is, cases of variability that involve quite distinct categorical realizations. For example, this is the case with the allophones of /t/ and /d/ in English as documented by Zue and Laferriere (1979). There are also patterns of temporal/spatial gradience that are highly systematic, as numerous studies of coarticulation and phonetic implementation show.

These issues are also closely related to the question of sources of diachronic change and the issue of whether change is gradual. The nature of variation as it is manifested in the social system and its relationship to diachronic change are very important issues, but not ones that I pursue here. (See work by Labov, Scobbie, Bybee, Kiparsky for recent discussions.)

**2.2.1.3 Gradient well-formedness** There is gradience across the lexicon, or statistical knowledge, as documented in recent work by Pierrehumbert, Frisch, and others. (See Frisch 2000 for a review and Bod *et al.* 2003 for recent discussion.) Here we talk about gradient well-formedness, the idea that speaker/hearers make relative judgements about the well-formedness of various sound structures. In the case of phonotactics, this is understood as resulting from stochastic generalizations across the lexicon. Such gradient well-formedness judgements are observed in other aspects of

the phonology, as well as other domains including both morphology and syntax. (See other chapters, this volume.) In such cases, it is the judgement about well-formedness or grammaticality that is gradient, not a physical event in time and space such as in the first sense.

We turn now to the question of how gradience might be manifested in the different facets of phonology, focusing primarily on temporal/spatial gradience and gradient well-formedness.

### 2.2.2 Contrast

Fundamental to a phonological system is the idea of lexical contrast: some phonetic differences in the acoustic signal result in two distinct lexical items, that is, minimal pairs. This is also the basis upon which inventories of sounds are defined. The term *contrast* is used in two rather different senses: underlying or lexical contrast, and surface contrast, that is, identifiable phonetic differences independent of meaning. The question of surface contrast sometimes arises when comparisons are made between phonological categories across languages. It also often arises in the discussion of phonological alternations that affect lexical contrasts in terms of neutralization or near-neutralization. Cases of complete phonological neutralization should result in no cues to underlying differences or contrast. Yet many cases of what are claimed to be complete neutralization exhibit subtle phonetic cues that differentiate between surface forms. (For a recent discussion and review of such cases involving final devoicing, see Warner *et al.* 2004). Under one interpretation, such cases can be understood as gradient realization of contrast. Due to space limitations, I do not pursue the issue of near-neutralization here.

We might wonder if contrast is all or nothing, or whether it too might be gradient in the sense of exhibiting gradient well-formedness. Within generative grammar, we understand contrast in absolute terms. Two sounds are either in contrast or they are not. Many contrasts are very robust. Yet, contrast can also be much more specific or limited. (See Ladd 2003 for a discussion of some such cases.) There are certain sounds that contrast in some positions, but not others (that is, positional neutralization). For example, even for speakers who maintain an /a/ - /ɔ/ contrast in American English, this contrast holds only before coronals and in open syllables. What is the nature of realization of these sounds before non-coronals? Do speakers produce the 'same' vowel in *fog* and *frog*? There are also some sounds that contrast in all positions in the word, but where the functional load of the contrast is very limited, such as in the case of /θ/ versus /ð/ in English (*thigh vs. thy, ether vs. either, Beth vs. eth*, that is [ð]). Is contrast realized the same way in these cases

as in the more robust cases? Or should contrast also be understood as a gradient property? I will not pursue this question here, but it might well be that contrast is more gradient in nature than often assumed and so robustness of contrast might well prove to be an interesting area for investigation. Lexical neighbourhood effects as well as phonological conditioning might both come into play.

### 2.2.3 *Phonotactics*

A second aspect of sound systems widely understood to constitute part of phonology is allowable sound combinations or sequences—phonotactics. Some aspects of phonotactics appear to be defined by segmental context, especially immediately preceding and following elements; some aspects are defined by prosodic position, often best characterized in terms of syllable structure; and some aspects are defined by morpheme- or word-position. Under many approaches to phonology, phonotactic patterns are understood to be categorical in nature. Particular combinations of sounds are understood to be either well-formed or ill-formed. Following most generative approaches to phonology, both rule-based and constraint-based, phonotactic patterns are captured with the same formal mechanisms as phonological alternations. Typically, phonotactic and allophonic patterns closely parallel each other, providing the motivation for such unified treatments. It is argued that distinct treatments would result in a 'duplication' problem (e.g. Kenstowicz and Kisseberth 1977).

Recent work by a wide range of scholars (e.g. Pierrehumbert 1994, Vitevich *et al.* 1997, Frisch 2000, Bybee 2001, and Hay *et al.* 2003) suggests that phonotactic patterns can be gradient, in the sense that they do not always hold 100 per cent of the time. Phonotactic patterns may reflect the stochastic nature of the lexicon and speaker/hearers are able to make judgements about the relative well-formedness of phonotactic patterns.

As an example, consider the phonotactics of medial English clusters, as analysed by Pierrehumbert (1994). Pierrehumbert asks the question of how we can account for the distribution of medial clusters, that is, the fact that certain consonant sequences are well-formed but others are not, for example /mpr/, /ndr/ but not \*/rpm/ or \*/rdn/. A generative phonology approach predicts: medial clusters = possible codas + possible onsets. While a stochastic syllable grammar makes different predictions: 'the likelihood of medial clusters derived from the independent likelihoods of the component codas and onsets' (1994: 174) and 'The combination of a low-frequency coda and a low-frequency onset is expected to be a low-frequency occurrence' (1994: 169). Pierrehumbert carried out a systematic analysis of a dictionary and found

roughly fifty monomorphemic medial clusters. In the same dictionary, there were 147 possible codas and 129 possible onsets. If these were freely combining, there would be predicted to be 18,963 medial clusters. With some expected restrictions, Pierrehumbert concludes that we would still expect approximately 8,708. Pierrehumbert observes 'It turned out that almost all the occurring triconsonantal clusters were among the 200 most likely combinations, and that a stochastic interpretation of syllable grammar effectively ruled out a huge number of possible clusters, eliminating the need for many idiosyncratic constraints in the grammar' (1994: 169). Pierrehumbert then discusses the systematic restrictions that play a role in determining the particular fifty or so medial combinations that are attested among the 200 most likely. She concludes that a stochastic syllable grammar understood in the context of certain more traditional sorts of phonological constraints accounts for the observed patterns.

Recent work in psycholinguistics shows that speakers have access in at least some situations to very fine details including both speaker-specific and situation-specific information. (See Beckman 2003 and Pierrehumbert 2003 for reviews and discussion of this body of work.) Thus, it is not that surprising that speakers are sensitive to degrees of well-formedness in phonotactic patterns and that these parallel in some cases distributions in the lexicon.

This leads us to two important issues. First, are phonotactic patterns and other aspects of phonology (contrast, morphophonemics, and allophony) as closely associated with each other as has been assumed in the generative phonological literature? Perhaps while similar and in some cases overlapping, phonotactics and other aspects of phonological patterning are not necessarily the same thing. This suggests that the standard generative phonology approach is reductionist in that it collapses distributional generalizations across the lexicon with other aspects of what is understood to be phonology. Second, evidence suggests that we have access to finer details in at least some situations/tasks and some of these finer details may play a role in characterizing lexical entries. Thus, it cannot be, as is often assumed following theories of underspecification in generative phonology, that lexical representations consist only of highly sparse contrastive information (e.g. *pit* /pit/, *spit* /spit/). We will not reach insightful conclusions about the nature of phonology if we just assume that lexical representations capture only contrast. These two widely held assumptions of generative phonology need to be revisited.

However, there are two important caveats on the other side. Just because we are sensitive to finer details, does not mean that we cannot abstract across the lexicon. To assume that we do not is to fall prey to this duplication problem from the other side. Pierrehumbert (2003) argues that some phonotactic

knowledge is not tied to frequency and indeed is true abstraction across the lexicon, that is, there is phonological knowledge independent of statistical generalizations across the lexicon. 'In light of such results, I will assume, following mainstream thought in linguistics, that an abstract phonological level is to be distinguished from the lexicon proper.' (2003: 191). This suggests that we have access to both fine-grained and coarse-grained levels of knowledge and that they co-exist (see Beckman 2003 and Beckman *et al.* 2004). We would predict a (negative) correlation between the degree of gradience and the level of abstraction.

#### 2.2.4 Alternations (morphophonemics)

In many ways, the core phenomena understood to constitute phonology are alternations. The most canonical types are morphophonemic alternations, where the surface form of a morpheme is systematically conditioned by phonological context. Alternation is also used to refer to allophonic alternation where particular phones are in complementary distribution and are thus argued to be variants of the same underlying phoneme. Positional allophones are argued to *alternate* in their distribution based on phonological context. We consider morphophonemic alternations in this subsection and allophony in Section 2.2.5.

Assuming we can draw appropriate boundaries (delineating the cases that are phonologically conditioned, productive, and not morpheme-specific), morphophonemic alternations are at the very core of what most phonologists think of as phonology. Most alternations are understood to be quite categorical in nature, often involving the substitution of distinct sounds in particular environments. Following a Lexical Phonology approach (e.g. Kiparsky 1982), such alternations are understood to be part of the lexical phonology and are assumed to respect structure preservation. If these sorts of cases are shown to involve gradience, this would strike at the core of our understanding of the phonology, since these are the least disputable candidates for 'being phonology'.

A widely cited claim arguing for gradience in phonology is that made by Steriade (2000). Parallel to phonological paradigm uniformity effects, which are taken to account for some 'cyclic' effects (e.g. Benua 1998; Kenstowicz 2002), Steriade argues that there are phonetic paradigm uniformity effects, where non-contrastive phonetic details may be marshalled to indicate morphological relatedness.

Consider first a canonical example of so-called paradigm uniformity effects. Many languages show *overapplication* or *underapplication* of phonological patterns that result in phonological similarity within morphologically related forms, despite the absence of the relevant phonological conditioning

context. For example, in Sundanese there is a general pattern of vowel nasalization, whereby vowels become nasalized after a nasal consonant, unless blocked by a non-nasal supra-laryngeal consonant (Robins 1957). This is exemplified in (2.2a). There is overapplication of nasalization in infixed forms indicating plurality or distributedness (2.2b).

#### (2.2) Nasalization in Sundanese (Cohn 1990)

- a. /ɲiar/ [ɲiār] 'seek' (active)  
 /niis/ [ni̯i̯s] 'relax in a cool place' (active)  
 /ɲatur/ [ɲātur] 'arrange' (active)  
 /ɲuliat/ [ɲūliat] 'stretch' (active)

- b. Singular Plural  
 /ɲiar/ [ɲiār] 'seek' (active) /ɲ=al=iar/ [ɲāliār]  
 /niis/ [ni̯i̯s] 'relax' (active) /n=ar=iis/ [nāri̯i̯s]

In derivational approaches, this overapplication follows from a cyclic analysis, where vowel nasalization reapplies after infixation (e.g. Cohn 1990). However such a solution is not available in non-serial approaches such as most Optimality Theoretic approaches. One account within Optimality Theory is that such patterns result from Output-Output constraints, reflecting the morphological relationships between words (Benua 1998). Such phonological parallels are enforced by *paradigm uniformity*.

Steriade (2000) argues that not only phonological properties (those that are potentially contrastive) show such effects but that 'paradigmatic uniformity is enforced through conditions that govern both phonological features and properties presently classified as phonetic detail, such as non-contrastive degrees in the duration of consonant constrictions, non-contrastive details in the implementation of the voicing contrast, and degree of gestural overlap.' (2000: 314). She then goes on to say that 'There is a larger agenda behind this argument: the distinction between phonetic and phonological features is not conducive to progress and cannot be coherently enforced.' (2000: 314)

This very strong claim rests on two cases. The first case is schwa deletion in French, where paradigm uniformity is argued to be responsible for the subtle differences between forms such as *pas d'rolé* 'no role' and *pas drôle* 'not funny', where the syllable-initial character of [ʁ] is maintained in the first case, despite the deletion of schwa. The second is flapping in American English: the observation (made by Withgott 1983 and others) that in some cases where the phonological environment is met for flapping, flapping does not occur is argued to be due to subphonemic paradigm uniformity, for example *cápiDalist*: *cápiDal*, but *militaristic*: *military*.

In Steriade's argument concerning flapping there are two crucial assumptions. First, 'we suggest that PU [paradigm uniformity] (STRESS) should characterize not only stress identity between syllables but also the use of individual stress correlates (such as duration, pitch accents, vowel quality) to flag the stress profile of the lexical item to be accessed.' (2000: 321). In effect, what Steriade hopes to conclude—that non-contrastive details can drive paradigm uniformity—becomes a working assumption, making the argument circular. Second, 'The difference between [r] and [t]/[d] is a function of closure duration.... The extra-short duration of [r] is a candidate for a never-contrastive property' (2000: 322). In fact, there are a number of other candidates for the difference between flap and [d/t], some of which are contrastive properties (such as sonority). Steriade conducted an acoustic study of twelve speakers uttering one repetition each of several pairs of words, with judgements based on impressionistic listening (which turns out to be rather unreliable in identifying flapping). Based on the results of the study, she concludes that PU (STRESS: DURATION) is responsible for observed base-derivative correspondence.

However a recent experiment designed to replicate Steriade's finding by Riehl (2003a, 2003b) calls into question Steriade's (2000) conclusions about the nature of the paradigm uniformity effect. Riehl recorded six speakers, with twelve repetitions of each form, using similar pairs to those in Steriade's study. She undertook an acoustic analysis of the data (including measures of closure duration, VOT, presence or absence of burst, and voicing duration during closure) and also a systematic perceptual classification by three listeners, in order to compare the perception and consistency of perception with the acoustic realization.

In Riehl's data, there were four relevant pairs of forms where phonologically one might expect a flap in one case and a stop in the other, as in the examples studied by Steriade. There were 24 possible cases of paradigm uniformity (4 forms × 6 speakers) where 12/12 forms could have shown both flaps or both stops. Since there was quite a bit of variation in the data, Riehl counted either 12/12 or 11/12 cases with the same allophone as showing 'uniformity'. Out of the 24 cases, there were 7 that showed uniformity or near uniformity and 17 with variation within forms and within pairs. Thus the case for paradigm uniformity was weak at best. In cases of variation, stops were usually produced earlier in the recordings, flaps later, arguably showing a shift from more formal to more casual speech (highlighting the importance of looking at multiple repetitions). Moreover, Riehl found that the coding of particular tokens as flaps or stops was not as straightforward as often assumed and she found that the perception of flaps correlated best with VOT, not closure duration.

This does not mean that there is no morphological influence on flapping, but suggests that the pattern may not be that strong. There is also a lack of compelling evidence to show that these effects are best understood as sub-phonemic paradigm uniformity. Steriade's conclusions regarding French schwa are also not that secure. It is not clear whether these effects are really what we understand to be paradigm uniformity; rather, this interpretation seems to be driven by Steriade's assumption that phonology and phonetics are not distinct. (Barnes and Kavitskaya 2002 also question Steriade's conclusions in the case of schwa deletion in French.) Does this mean that there are not gradient effects in the domain of morphophonemics? A more convincing case of morphological influences on phonetic realization may be the degree of glottalization in English correlating with degree of morphological decompositionality, for example, *realign* versus *realize* as discussed by Pierrehumbert (2002). 'The model predicts in particular the existence of cases in which relationship of phonetic outcomes to morphological relatedness is gradient.' (2002: 132) The question is how close the correlation between morphological decompositionality and phonetic realization is and how best to model this correlation. I fully agree with Pierrehumbert that 'More large-scale experiments are needed to evaluate this prediction.' (2002: 132)

### 2.2.5 Allophony

The final aspect of phonology is allophony. Based on the definitions of *SPE*, allophony is understood to be part of phonology, due to its language-specific nature. There has been much discussion in the literature about whether allophony is necessarily categorical in nature or whether there are gradient aspects of allophony. There are also many cases of what was understood as allophony in categorical terms that have been shown, based on instrumental studies, to be gradient. This is the case of anticipatory nasalization in English discussed in Cohn (1990, 1993) and the case of velarization of [l] in English as discussed by Sproat and Fujimura (1993). Such cases raise three issues.

1. Based on impressionist description, much work on allophony suggests that allophony is quite categorical in nature. Yet, both the tools we use (careful listening) and the symbols available to us (phonetic transcription which is discrete in nature) bias our understanding of these patterns as categorical.
2. There has been a wide body of work arguing for a rethinking of the *SPE* definition of what is phonology and what is phonetics. Much work has identified the language-specific nature of phonetic patterns (e.g. Chen 1970; Keating 1985; Cohn 1990; Kingston and Diehl 1994), leading to a



rethinking of where we draw the boundary between phonetics and phonology. Under these approaches many cases that have been thought of as phonological have been reanalysed as phonetic.

3. This still leaves us with the question of where to draw the line and whether we should draw a line. We return to this question in Section 2.3.

There has also been argued to be gradience in allophony in a rather different sense. Bybee (2001) and Jurafsky *et al.* (2001) among others argue that lexical (token) frequency affects allophony in the sense that more frequent words are observed to be shorter and phonologically reduced. Bybee (2001 and earlier work) has argued that what is understood as allophony in generative phonology cannot follow from general rules or constraints, because there are frequency effects on the realization of non-contrastive properties. If what we think is allophony falls along a continuum rather than in two or three discrete categories, and if there is a strong correlation between the realization of a particular non-contrastive property and frequencies of particular lexical items in the lexicon, then this would be difficult to model in standard generative phonological models.

One widely cited case in this regard is the case of schwa deletion in the context of a resonant in English. Bybee (2001) citing an earlier study based on speaker self-characterization (Hooper 1976, 1978, a.k.a. Bybee) argues that it is not the case that schwa is either deleted or present, but rather that there are degrees of shortening. She observes impressionistically a continuum from [Ø], to syllabic resonants, to schwa + resonant, for example *every* [Ø], *memory* [r], *mammary* [ər] (where a syllabic resonant is thought to be shorter than a schwa plus resonant). It is further argued that these different realizations correlate with the lexical (token) frequency, that is, there is complete deletion in the most common forms, schwa plus a resonant in the least frequent forms, and syllabic resonants in the cases which fall in between. This is understood to follow from the view that sound change is lexically and phonetically gradual, so that 'schwa deletion' is farther along in high frequency words.

Lavoie (1996) tried to replicate Bybee's finding with a more systematic study including instrumental data. Her study included acoustic measurements of multiple repetitions of near minimal triplets, sets that were similar in their phonological structure and differed in relative frequency both within the sets and in terms of absolute frequency across the data set, based on frequency from Francis *et al.* (1982). Crucially when frequency was plotted against duration, no correlation was found. Rather, there was a robust subpattern of complete deletion of schwa in many forms independent of

lexical frequency and there was variation in duration independent of lexical frequency. Thus schwa deletion in English does not provide the kind of evidence that Bybee suggests for allophony being driven by lexical token frequency. (The other cases widely discussed by Bybee in this regard, such as aspiration of /s/ in Spanish and /t, d/ deletion in English also warrant careful reconsideration.)

We need to consider the question of whether there are cases where gradient phonological patterns correlate with lexical (token) frequency. The short answer is yes, but the best documented cases in this regard are of a very different sort than those mentioned above. When token frequency differences correlate with function versus content word differences, indeed frequency in such cases has a major effect on the realization of sound patterns. Function words show much more reduced and variable realization than content words. See for example Lavoie (2002)'s study of the reduction and phonetic realization of *for* versus *four* and Jurafsky *et al.*'s (2001) study of both reduction and variability in function words. Bybee assumes that it is token frequency that differentiates function words and content words; yet these frequency effects can also be understood to follow from the prosodic difference between content and function words. (For recent discussion of the prosodic structure of lexical versus functional categories, see e.g. Zec 2002.)

Unequivocal support for Bybee's claim would come from duration differences strongly correlated with token frequency differences found within the same lexical category, with appropriate controls for discourse context, priming effects, and so forth. Cohn *et al.* (2005) investigate the phonetic durations of heterographic pairs of homophonous English nouns that differ in token frequency. Homophonous pairs were grouped into three categories based on the magnitude of the frequency difference between the members of each pair, as determined by relative frequencies in five large corpora. This included Large Difference pairs (e.g. *time* ~ *thyme*, *way* ~ *wey*), Medium Difference pairs (e.g. *pain* ~ *pane*, *gate* ~ *gait*), and little or No Difference pairs (e.g. *son* ~ *sun*, *peace* ~ *piece*). Four native speakers of American English participated in two experiments. In the first experiment, the speakers were recorded reading four repetitions of a randomized list of the target words in a frame sentence. In the second experiment, a subset of these words was read in composed sentences with controlled prosodic structures. The phonetic duration of each target word was then measured in Praat, and the ratio more frequent/less frequent was calculated for each repetition of each pair. If the hypothesis that greater frequency leads to shorter duration is correct, then these ratios should systematically fall below one for the Large Difference and Medium Difference pairs, while those for the little or no difference group should be approximately

one. No systematic differences were found for individual speakers or across speakers in either the frame sentences or the composed sentences. The lack of positive correlation between duration and token frequency calls into question the hypothesis that greater frequency leads to shorter duration. These results are interesting in light of Jurafsky's (2003) observation that evidence for frequency effects are better documented in comprehension than production. On the production side, effects are much more robustly documented for latency in lexical access than in phonetic duration differences. These results and observations highlight the need for a better understanding of the locus of frequency effects in the lexicon and in speech production.

### 2.3 Conclusions and implications

Having considered the evidence for three cases of gradience in the phonology in Sections 2.2.3–5, we now return to the broader question: Is there gradience in the phonology? Not surprisingly, the answer seems to be yes and no. It depends on what we mean by gradience and it depends on which facets of the phonology we consider. The clearest evidence for gradience among the cases that we considered is gradient well-formedness, as documented in the case of phonotactics. It was less clear that there was a convincing empirical basis for the specific claims made by Steriade in terms of subphonemic effects in paradigm uniformity and those made by Bybee regarding frequency effects in allophony. However the shakiness of the specific cases does not answer the question of whether there is gradience in phonology in the areas of morphophonology and allophony. In both cases, the conclusion about whether there is gradience in the phonology depends in part on the definition of phonology and how we understand phonology in relationship to phonetics.

This leads us back to the question, discussed in Section 2.1.2, whether phonetics and phonology are distinct domains. A modular view of the grammar necessarily leads us to a mapping approach between the phonology and phonetics view. On the other hand, focusing primarily on the grey area, cases that are particularly difficult to classify, and defining similarity as 'duplication' lead us to a unidimensional view.

Let us return to the observation by Pierrehumbert *et al.* (2000) that knowledge of sound structure falls along a continuum, with more fine-grained knowledge tending to lie at the phonetic end and lexical contrast and alternations being more granular. This continuum is schematized in Figure 2.2a with phonetics versus phonology on the x-axis and degree of granularity on the y-axis. Consider the schematic distribution of the data: the modular approach suggests a distribution such as that in Figure 2.2b, with




FIGURE 2.2. (a) Continuum between phonetics and phonology (x-axis) and fine-grained and granular (y-axis) dimensions of speech; (b) distribution of data, modular approach; (c) distribution of data, unidimensional approach

little or no grey area. The unidimensional approach suggests a distribution such as that in Figure 2.2c, with little correlation between the two dimensions. Yet the data clearly fall somewhere between these two views. How can we understand and model this distribution?

Two methodological issues contribute to the perceived cost of 'duplication' and to the tendency to avoid duplication through reductionism. The first is the nature of modularity. Hale and Reiss (2000: 162) state 'The modular approach to linguistics, and to science in general, requires that we both model the interactions between related domains, and also sharply delineate one domain from another.' But, we need to ask the question: Is there strict modularity? Does modularity entail sharp delineation? Could there be modularity that is not rigid? The lack of strict modularity is implicit in views to understanding the relationships between linguistic domains through interfaces. If we do not subscribe to strict modularity between phonology and phonetics and between phonology and the lexicon, then it becomes an empirical question if drawing a distinction is useful. Does a division of labour contribute to both descriptive adequacy and explanatory adequacy?

The second is the status of Occam's Razor, or the principle of parsimony. Perhaps Occam's Razor does not hold as strongly as we believe. There is redundancy in language. Redundancy is widely observed in the domain of phonetics in terms of multiple and varied cues to the realization of particular phonological structures. Even cases of what we understand to be

a straightforward phonological contrast may involve multiple cues. Evidence suggests that lexical representations include multiple levels of details, including the kind of sparse abstract representations widely assumed in generative phonology and much more fine-grained levels of detail. (See Beckman *et al.* 2004 for discussion and a specific proposal in this regard.) Not only is there redundancy within domains, but there appears to be redundancy across domains, so 'duplication' is not a problem, but in fact an intrinsic characteristic of language. Increasingly there is agreement that unidimensional or reductionist views are not sufficient (see Pierrehumbert 2001: 196). Attempting to understand sound structure in only abstract categorical terms or in only the gradient details, or trying to understand the nature of the lexicon in exactly the same terms in which we try to understand phonology is insufficient.

In conclusion, the relationship between phonetics and phonology is a multifaceted one. It reflects phonetic constraints that have shaped synchronic phonological systems through historical change over time. Synchronically, phonological systems emerge as a balance between the various demands placed on the system, but the evidence suggests that phonology cannot be reduced to the sum of these influences. Phonetics and phonology also need to be understood in relationship to the lexicon. There are parallels and overlaps between these three areas, but none of them is properly reduced to or contained in the others. Language patterns are fundamentally fluid. There is evidence of phonologization, grammaticalization, lexicalization, and so forth. Similar patterns can be observed across these domains. To reach a fuller understanding of the workings of the sound system and the lexicon, we need to be willing to reconsider widely held assumptions and ask in an empirically based way what is the connection between these domains of the linguistic system.

## 3

## Gradedness: Interpretive Dependencies and Beyond

ERIC REULAND

### 3.1 Introduction

During the last decades it has been a recurrent theme whether or not the dichotomy between grammatical versus ungrammatical or well-formed versus ill-formed should not be better understood as a gradient property (cf. Chomsky's (1965) discussion of degrees of grammaticality).<sup>1</sup> If so, one may well ask whether gradedness is not an even more fundamental property of linguistic notions. The following statement in the announcement of the conference from which this book originated presupposes an affirmative answer, and extends it to linguistic objects themselves, making the suitability to account for gradedness into a test for linguistic theories: 'The kind of grammar typically employed in theoretical linguistics is not particularly suited to cope with a widespread property of linguistic objects: gradedness.'<sup>2</sup> This statement implies that we should strive for theories that capture gradedness. To evaluate it one must address the question of what 'gradedness' as a property of linguistic objects really is. The issue is important. But it is also susceptible to misunderstandings. My first goal will be to show that gradedness is not a unified phenomenon. Some of its manifestations pertain to language use rather than to grammar *per se*. Understanding gradedness may therefore help us shed light on the division of labour among the systems underlying language and its use. Showing that this is the case will be the second goal of this contribution.

<sup>1</sup> This material was presented at the 'Gradedness conference' organized at the University of Potsdam 21–23 October 2002. I am very grateful to the organizers, in particular Gisbert Fanselow, for creating such a stimulating event. I would like to thank the audience and the two reviewers of the written version for their very helpful comments. Of course, I am responsible for any remaining mistakes.

<sup>2</sup> This statement is taken from the material distributed at the conference.

Is there gradient phonology?, A Cohn (2006)

## References

- Antilla, A. (2002) Variation and phonological theory. In J. Chambers, P. Trudgill, and N. Schilling-Estes (eds) *The Handbook of Language Variation and Change*. Oxford: Blackwell, pp. 206-243.
- Barnes, J. and D. Kavitskaya (2002) Phonetic analogy and schwa deletion in French, presented at the Berkeley Linguistic Society.
- Beckman, M. (2003) Input representations (inside the mind and out). *WCCFL 22 Proceedings*, ed. M. Tsujimura and G. Garding. Somerville, MA: Cascadia Press, pp. 101-125.
- Beckman, M., B. Munson, and J. Edwards (2004) Vocabulary growth and developmental expansion of types of phonological knowledge. LabPhon 9 preconference draft.
- Benua, L. (1998) *Transderivational Identity*. UMass PhD dissertation.
- Blevins, J. (2004) *Evolutionary Phonology: The Emergence of Sound Patterns*. Cambridge: Cambridge University Press.
- Bod, R., J. Hay and S. Jannedy (2003) *Probabilistic Linguistics*. Cambridge, MA: The MIT Press.
- Browman, C. and L. Goldstein (1992) Articulatory Phonology: an overview. *Phonetica* 49: 155-180.
- Burton-Roberts, N., Carr, P. & G. Docherty (2000) *Phonological Knowledge: Conceptual and Empirical Issues*. New York: Oxford University Press.
- Bybee, J. (2001) *Phonology and Language Use*. Cambridge: CUP.
- Chen, M. (1970) Vowel length variation as a function of the voicing of consonant environment. *Phonetica* 22: 129-159.
- Chomsky, N. and M. Halle (1968) *The Sound Pattern of English*. New York: Harper & Row.
- Clements, G. N. (1992) Phonological primes: Gestures or features? *Working Papers of the Cornell Phonetics Laboratory* 7: 1-15.
- Coetzee, A. (2004) *What it Means to be a Loser: Non-optimal Candidates in Optimality Theory*. UMass PhD dissertation.
- Cohn, A. (1990) *Phonetic and Phonological Rules of Nasalization*. PhD dissertation, UCLA. Distributed as UCLA Working Papers in Phonetics 76.
- Cohn, A. (1993) Nasalization in English: Phonology or phonetic. *Phonology* 10: 43-81.
- Cohn, A. (1998) The phonetics-phonology interface revisited: Where's phonetics? *Texas Linguistic Forum* 41: 25-40.
- Cohn, A. (2003) Phonetics in phonology and phonology in phonetics. Paper presented at 11<sup>th</sup> Manchester Phonology Meeting, Manchester, UK.
- Cohn, A., J. Brugman, C. Clifford, and A. Joseph (2005) Phonetic Duration of English Homophones: An Investigation of Lexical Frequency Effects, presented at LSA, 79<sup>th</sup> meeting, Oakland, CA.
- Flemming, E. (2001) Scalar and Categorical Phenomena in a Unified Model of Phonetics and Phonology. *Phonology* 18: 7-44.
- Francis, W. Nelson and Henry Kucera (1982) *Frequency Analysis of English Usage: Lexicon and Grammar*. Boston: Houghton Mifflin.
- Frisch, S. (2000) Temporally organized lexical representations as phonological units. In M. Broe and J. Pierrehumbert (eds.) *Papers in Laboratory Phonology V: Acquisition and the Lexicon*. Cambridge, England: Cambridge University Press, pp. 283-298.
- Hale, M. and C. Reiss (2000) Phonology as cognition. In Burton-Roberts, N., Carr, P. & G. Docherty (eds.) *Phonological Knowledge: Conceptual and Empirical Issues*. New York: Oxford University Press. pp. 161-184.
- Hay, J. Pierrehumbert, J. and M. Beckman (2003) Speech perception, wellformedness, and the statistics of the lexicon. In J. Local, R. P. Gden, and R. Temple (eds.) *Phonetic Interpretation: Papers in Laboratory Phonology VI*. Cambridge: CUP.
- Hayes, B., R. Kirchner, and D. Steriade (2004) *Phonetically Based Phonology*. Cambridge: CUP.
- Hooper, Joan B. (1976) "Word Frequency in Lexical Diffusion and the Source of Morphophonological Change." In Christie, W. (ed.) *Current Progress in Historical Linguistics*. Amsterdam: North Holland, 95-105.
- Hooper, Joan B. (1978) "Constraints on Schwa Deletion in American English." In Fisiak, J. (ed) *Recent Developments in Historical Phonology*. The Hague: Mouton, 183-207.
- Hume, E. and K. Johnson (2001) *The Role of Speech Perception in Phonology*. San Diego: Academic Press.
- Hyman, L. (1976) Phonologization. In A. Juillard (ed.) *Linguistic Studies Offered to Joseph Greenberg*. Vol. 2 Saratoga: Anna Libri, pp. 407-418.
- Jurafsky, D. (2003) Probabilistic modeling in psycholinguistics: Linguistic comprehension and production. In R. Bod, J. Hay and S. Jannedy (eds.) *Probabilistic Linguistics*. Cambridge, MA: The MIT Press, pp. 39-95.
- Jurafsky, D. Bell, Gregory and Raymond (2001) Probabilistic relations between words: Evidence from reduction in Lexical Production, in Bybee and Hopper (eds) *Frequency and the emergence of linguistic structure*. Amsterdam: John Benjamins.
- Keating, P. (1985) Universal phonetics and the organization of grammars. In V. Fromkin (ed) *Phonetic Linguistics: Essays in Honor of Peter Ladefoged*. Orlando: Academic Press, pp. 115-132.

- Keating, P. (1988) The window model of coarticulation: articulatory evidence. *UCLA Working Papers in Phonetics* 69: 3-29.
- Keating, P. (1996) The Phonology-Phonetics interface. *UCLA Working Papers in Phonetics* 92: 45-60.
- Kenstowicz, M. (2002) Paradigmatic uniformity and contrast. *MIT Working Papers in Linguistics* vol. 42 Phonological Answers (and Their Corresponding Questions), pp#?
- Kenstowicz, M and C. Kisseberth (1977) *Topics in Phonological Theory*. New York: Academic Press.
- Kingston, J. and R. Diehl (1994) Phonetic Knowledge. *Language* 70: 419-454.
- Kiparsky, P. (1982). Lexical morphology and phonology. In The Linguistics Society of Korea (eds.) *Linguistics in the Morning Calm*. Seoul: Hanshin Publishing Co. pp. 3-91.
- Kirchner R. (2001) *An Effort-Based Approach to Consonant Lenition*. New York, NY: Routledge. [1998 UCLA Ph.D dissertation].
- Klatt, D. (1987) Review of text-to-speech conversion for English. *JASA* 82.3: 737-793.
- Ladd, D. R. (2003) "Distinctive phones" in surface representation. Written version of paper presented at LabPhon 8, to appear in the Proceedings.
- Lavoie, L. (1996) Lexical frequency effects on the duration of schwa—Resonant sequences in American English. Poster presented at LabPhon 5, Chicago, IL June 1996.
- Lavoie L. (2002) Some influences on the realization of *for* and *four* in American English. *JIPA* 32: 175-202.
- McCarthy, J. (2003) OT constraints are categorical. *Phonology* 20: 75-138.
- McCarthy, J. and A. Prince (1993) Generalized alignment. In *Morphology Yearbook 1993*. G. Booij and J. van Marle (eds.). Dordrecht: Kluwer, pp. 79-153.
- Pierrehumbert, J. (1980) *The Phonology and Phonetics of English Intonation*. MIT Ph.D. dissertation.
- Pierrehumbert, J. (1994) Syllable structure and word structure: A study of triconsonantal clusters in English. In Keating (ed.) *Phonological Structure and Phonetic Form: Papers in Laboratory Phonology III*, Cambridge: CUP, pp. 168-188.
- Pierrehumbert, J. (2001) Stochastic phonology. *GLOT* 5 No.6: 195-207.
- Pierrehumbert, J. (2002) Word-specific phonetics. In Gussenhoven and Warner (eds) *Laboratory Phonology 7*. Mouton de Gruyter, pp. 101-139.
- Pierrehumbert, J. (2003) Probabilistic Phonology: Discrimination and Robustness. In R. Bod, J. Hay and S. Jannedy (eds.) *Probabilistic Linguistics*. Cambridge, MA: The MIT Press, pp. 177-228.
- Pierrehumbert, J., Beckman, M. E and Ladd, D. R. (2000) Conceptual foundations in phonology as a laboratory science. In Burton-Roberts, N., Carr, P. and Docherty, G. (eds) *Phonological Knowledge: Conceptual and Empirical Issues*, New York: Oxford University Press, 273-304.
- Riehl, A. (2003a) American English flapping: Evidence against paradigm uniformity with phonetic features. *Proceedings of the 15th International Congress of Phonetic Sciences*, pp. 2753-2756.
- Riehl, A. (2003b) American English flapping: Perceptual and acoustic evidence against paradigm uniformity with phonetic features. *Working Papers of the Cornell Phonetics Laboratory* 15. 271-337.
- Robins, R. H. (1957). Vowel nasality in Sundanese: a phonological and grammatical study. In *Studies in Linguistics* (special volume of the Philological Society). Oxford: Basil Blackwell, pp. 87-103.
- Scobbie, J. (2004) Flexibility in the Face of Incompatible English VOT Systems. Written version of *Lab Phon 8* paper.
- Sproat, R. and O. Fujimura (1993) Allophonic variation in English /l/ and its implications for phonetic implementation. *Journal of Phonetics* 21: 291-311.
- Steriade, D. (1990) Gestures and autosegments: Comments on Browman and Goldstein's paper. In Kingston and Beckman (eds.). 382-397.
- Steriade, D. (2000) Paradigm uniformity and the phonetics-phonology boundary, in M. Broe and J. Pierrehumbert (eds.) *Papers in Laboratory Phonology V: Acquisition and the Lexicon*. Cambridge, England: Cambridge University Press, pp. 313-334.
- Steriade, D. (2001) Directional asymmetries in assimilation: A directional account. In E. Hume and K. Johnson (eds.) *The Role of Speech Perception in Phonology*. San Diego: Academic Press, pp. 219-250.
- Vitevitch, M., P Luce, J. Charles-Luce, and D. Kemmerer (1997) Phonotactics and syllable stress: Implications for the processing of spoken nonsense words. *Language and Speech* 40: 47-62.
- Warner, N. A. Jongman, J. Sereno, and R. Kemps (2004) Incomplete Neutralization and Other Sub-Phonemic Durational Differences in Production and Perception: Evidence from Dutch. *Journal of Phonetics* 32: 251-276.
- Withgott, M. (1983) *Segmental Evidence for Phonological Constituents*. PhD dissertation, University of Texas, Austin.
- Zec, D. (2002) On the prosodic status of function words. *Working Papers of the Cornell Phonetics Laboratory* 14: 206-248.
- Zsiga, E. (2000) Phonetic Alignment Constraints: Consonant Overlap and Palatalization in English and Russian. *Journal of Phonetics* 28: 69-102.
- Zue, V. and M. Laferriere (1979) Acoustic Study of Medial /t, d/ in American English. *JASA* 66: 1039-1050.