9

Determining Cross-Linguistic Phonological Similarity Between Segments

The Primacy of Abstract Aspects of Similarity

Charles B. Chang

1 Introduction

The notion of *phonological similarity* – that is, similarity between two sound structures – is central to much of the literature on spoken language. Phonological similarity is invoked to explain a variety of systematic patterns in word recall (e.g., Copeland and Radvansky 2001; Fournet *et al.* 2003), lexical and conceptual development (e.g., Sloutsky and Fisher 2012), language games (e.g., Zwicky and Zwicky 1986), first-language (L1) and second-language (L2) perception (e.g., Johnson 2003; Best and Tyler 2007), L1 and L2 production (e.g., Major 1987; Page *et al.* 2007), loanword phonology (e.g., Kang 2003, 2008), and cross-linguistic interaction in bilingualism (e.g., Flege 1995; Laeufer 1996).

Several different kinds of "phonological similarity" are referred to in the literature, however, and these various types of similarity have diverse consequences for grammar and learning (for a recent overview, see Gallagher and Graff 2012). For example, some studies examine the effects of phonological similarity between lexical items – operationalized as "neighborhood density" – on speech perception and production (e.g., Luce and Pisoni 1998; Vitevitch 2002; Munson and Solomon 2004; Gahl *et al.* 2012), while other studies consider the similarity between the various potential forms of a lexical item in explaining distributional regularities such as phonotactic restrictions and environments for alternation and neutralization (e.g., Pierrehumbert 1993; Flemming 2004; Frisch *et al.* 2004; Steriade 2009; Gallagher 2012). Phonological similarity between individual sounds or natural classes of sounds has been measured perceptually via perceptual confusions or explicit mappings with goodness-of-fit ratings (e.g., Miller and Nicely 1955; Strange 1999; Best *et al.* 2003; Chang 2009b). Computational work, on the other hand, draws on a feature-based

The Segment in Phonetics and Phonology, First Edition. Edited by Eric Raimy and Charles E. Cairns.

type of phonological similarity to align segmental sequences, whether for the purposes of analyzing cognate relationships or developmental speech patterns (e.g., Covington 1996; Kondrak 2003; Kessler 2005). Importantly, a similarity metric that provides a good model of behavior in one case may make poor predictions in another. As noted by Gallagher and Graff (2012), perception and production data do not necessarily converge on the same similarity relations, nor do phonetic and phonological data (Mielke 2012).

The mismatch between "phonetic" kinds of similarity and "phonological" kinds of similarity is at the heart of a disparity that is commonly seen between segmental similarity relations within one language and those between two languages. In this chapter, I describe this mismatch in more detail and argue that conflicts between different types of similarity are so often resolved in the same way (namely, in favor of "phonological" kinds of similarity) because high-level information is weighed more heavily than low-level information. Note that the segment is fundamental to this argument because there is no clear way to implement the phonemic-level interactions described in this chapter without positing an abstract, segment-sized category such as the phoneme. In a sense, then, the cross-linguistic phenomena reviewed here can be considered evidence for the existence of segments as discrete phonological units, as well as for the very distinction between phonetics and phonology (in particular, their hierarchical relationship).[1]

The chapter is organized as follows. In section 2, I decompose the construct of phonological similarity into subtypes of similarity and review the problem of conflict between different types of similarity observed in the cross-linguistic speech literature. In section 3, I present an array of findings from cross-linguistic research showing a preference for relating segments and natural classes to each other on an abstract level. In section 4, I discuss the implications of such abstract knowledge for studies of cross-linguistic phonetics and phonology, and in section 5, I provide concluding remarks.

2 Components of phonological similarity and their interaction

2.1 Levels of similarity

The construct of phonological similarity can be decomposed into at least[2] three subtypes of similarity: objective acoustic similarity, language-specific allophonic similarity, and cross-linguistic phonemic similarity. These metrics of similarity have analogues in other models of phonological similarity that distinguish between various factors influencing overall similarity (e.g., Austin 1957; Ladefoged 1969; Flege 1996; Bohn 2002). Let us consider each type of similarity in turn.

Acoustic similarity refers to the raw (i.e., non-language-specific) distance between sounds in terms of acoustic dimensions such as frequency, duration, and amplitude. At a basic auditory level, listeners tend to perceive sounds that are relatively close acoustically (e.g., [f] and [θ]; [i] and [ɪ]) as more similar than sounds that are relatively distant acoustically (e.g., [s] and [θ]; [i] and [a]). For example, in a speeded discrimination task (thought to reflect non-linguistic perception of auditory contrast), native Dutch speakers and native English speakers take a comparably longer amount of time to discriminate the acoustically similar [f] and [θ] than the acoustically dissimilar [s] and [θ], even though these pairs of sounds are contrastive in English only (Johnson and Babel, 2010). This kind of result is consistent with the view that there is an acoustic/auditory basis for perceived similarity that transcends linguistic knowledge.[3] Nevertheless, differences in language background are likely to result in divergent perceptual patterns in linguistic tasks due to effects of allophonic similarity.

Allophonic similarity is based on within-language comparisons between sounds at the level of contextually defined allophones, which are specific to a particular language. A pair of sounds is allophonically similar to the extent that they can be related to each other within a language – by virtue of the fact that they do not contrast and/or the fact that they alternate with each other in a productive pattern (Johnson and Babel 2010). Consequently, a pair of sounds can be perceived differently by listeners of different language backgrounds if the two sounds exist in an allophonic relationship in one language, but not the other. For example, English speakers (for whom [d] contrasts with [ð] and alternates with [ɾ]) perceive [d] as more similar to [ɾ] than to [ð]; in contrast, Spanish speakers (for whom [d] contrasts with [ɾ] and alternates with [ð]) perceive [d] as more similar to [ð] than to [ɾ] (Boomershine *et al.* 2008). Similar patterns are found when a sound is absent from one language, but present in another. For instance, Dutch speakers (for whom [θ] does not occur as a phoneme) rate [θ] as more similar to [s] and [ʃ] than do English speakers (for whom all three fricatives are phonemes) (Johnson and Babel 2010).

While acoustic similarity is not specific to any language and allophonic similarity is specific to one language, phonemic similarity is related to sounds in two languages. Therefore, phonemic similarity is inherently cross-linguistic. Phonemic similarity is also abstract, because it is based on cross-language comparisons between sounds at the level of context-free phonemes, which may be viewed in terms of feature bundles. A pair of sounds that are acoustically and/or allophonically dissimilar may nonetheless be phonemically similar due to at least two factors: (1) similar positions in the respective phonemic inventories (when considering the contrastive feature oppositions – or, more broadly, the "relative phonetics"[4] – of the sounds in relation to other sounds in the inventory), and (2) similar distributional facts.

To take an example, American English and Mandarin Chinese both contain a vowel standardly transcribed as /u/ in their respective inventories, but the languages differ substantially in the quality of their /u/. English /u/ is acoustically far from Mandarin /u/ and much closer to the Mandarin front rounded vowel /y/ (Chang *et al.* 2011). Nevertheless, these two /u/ vowels can be identified as the "same" phoneme because they each occupy a similar place within the relevant vowel inventory – namely, that of a high back rounded vowel (i.e., [-consonantal, +syllabic, +high, +back, +round]). Even though English /u/ is relatively front and unrounded in comparison to Mandarin /u/, it is still the vowel that is the most high/back/rounded in the English inventory and, therefore, the most parallel to Mandarin /u/ in terms of vowel features. In addition to parallel inventory niches, English /u/ and Mandarin /u/ show similar distributional restrictions with the back rounded approximant /w/: neither can occur with /w/ in a stop-approximant onset cluster (i.e., *[pwu], *[twu], *[kwu], etc.). These similar co-occurrence restrictions suggest that, even though English /u/ is acoustically quite far from Mandarin /u/, they both pattern like back rounded vowels. In this way, English /u/ and Mandarin /u/ are phonemically similar despite their disparate phonetic realizations.

In summary, phonological similarity between segments can be said to exist at multiple levels: acoustic, allophonic, and phonemic. Acoustic similarity and allophonic similarity are relevant both within and between languages; however, phonemic similarity, since it involves the comparison of two phonological systems, is relevant only for cross-linguistic comparisons. As such, the perceived similarity between two segments within a language has typically been discussed in acoustic and/or auditory terms. In the next section, we review one influential attempt to encode this kind of perceptual similarity in the grammar and show how its predictions break down if extended to cross-linguistic comparisons.

2.2 Perceptual similarity in a (monolingual) grammar

Given how often linguistic phenomena are explained in terms of phonological similarity, it is reasonable to think that knowledge of similarity constitutes part of linguistic knowledge, and Steriade (2009) attempted to represent this knowledge in a language-universal "P-map," a set of ranked constraints regarding the relation of relatively similar vs. dissimilar forms. These constraints aim to maximize perceptual similarity between input and output forms, such that input-output correspondences between relatively dissimilar forms are penalized more heavily than those between relatively similar forms. For instance, in the case of final voiced stops, a typologically dispreferred structure, there might be two constraints – one penalizing a correspondence between a syllable-final voiced stop and a voiced stop-initial syllable containing an epenthetic vowel ($*D]_\sigma$–$DV]_\sigma$) and another penalizing a correspondence between a syllable-final voiced stop and a syllable-final voiceless stop ($*D]_\sigma$–$T]_\sigma$). Because a syllable-final voiced stop is arguably less similar to a new syllable than to a syllable-final voiceless stop (as reflected in similarity judgment data; Kawahara and Garvey 2010), the first constraint $*D]_\sigma$–$DV]_\sigma$ is ranked above the second constraint $*D]_\sigma$–$T]_\sigma$ ($*D]_\sigma$ –$DV]_\sigma >> *D]_\sigma$–$T]_\sigma$), such that, absent the influence of intervening constraints, syllable-final voiced stops are predicted to alternate with syllable-final voiceless stops, not with syllables containing an epenthetic vowel. In fact, this is what is found across a range of languages (e.g., Germanic and Slavic languages): final voiced stops are repaired via devoicing rather than epenthesis. This pattern is thus consistent with the basic prediction of the P-map – namely, that output patterns follow perceptual similarity relations between an input and its possible outputs.

But what about mapping between an L2 input and an L1 output? In a recent study, Shaw and Davidson (2011) showed that the tight link between perceptual similarity and ultimate output assumed by the P-map does not hold for cross-linguistic mapping. Controlling for a variety of factors, they observed that unfaithful production of novel (L2) input clusters cannot be said to follow from perceptual similarity, as fricative-stop clusters were produced with epenthetic forms (inserting a vowel in the middle of the cluster), despite being judged most similar to prothetic forms (inserting a vowel before the cluster). Some explanations offered for this unexpected disparity between production and perception were maximizing the perceptual recoverability of segments, as well as maintaining uniformity in repair strategy (given that stop-stop clusters were also produced with epenthetic forms). Still, it remains unclear why the P-map, which seems to do a good job accounting for within-language alternation, fails in this kind of cross-linguistic situation, since it is supposed to represent universal perceptual knowledge.

Here I consider the possibility that the P-map fails in cross-linguistic circumstances because, as a model of similarity based on within-language relations, it does not incorporate the influence of phonemic similarity between languages. The 'P' in the P-map stands for "perceptual," which reflects the fact that it encodes similarity relations based on (acoustic) perceptual similarity. However, as discussed above, acoustic similarity and allophonic similarity are not the only types of similarity that exist between L1 and L2 segments. Cross-linguistic phonological similarity may also be influenced by phonemic parallelisms, leaving the P-map ill-equipped to fully model cross-linguistic similarity relations. Thus, in the next section, we examine the hypothesis that phonemic similarity – in particular, its interaction with acoustic and allophonic similarity – plays a primary role in determining cross-linguistic similarity relations that depart from acoustic and allophonic comparisons (and, therefore, from predictions of the P-map).

2.3 Conflict and interaction between levels of similarity

The idea that phonemic similarity can result in cross-linguistic similarity relations differing from within-language similarity relations is based on two assumptions. The first assumption is that phonemic similarity sometimes differs from acoustic and allophonic similarity; as is discussed below in section 3, there is ample evidence that this situation actually arises. The second assumption is that at the phonetics-phonology interface there is a hierarchy between higher-level information (e.g., phonemic correspondence) and lower-level information (e.g., acoustic properties, allophonic alternations), with higher-level information taking precedence in cases of conflict.

In discussing this latter assumption, I proceed from an implication made by Flege (1996) in his definition of how to determine when a novel L2 sound is "new" vs. "similar" to an L1 sound. Flege observed that a useful heuristic in determining L1-L2 similarity is the "phonetic symbol test," one of Bohn's (2002) so-called "armchair methods": if an L1 sound and an L2 sound are transcribed with the same symbol in the International Phonetic Alphabet, this implies that the L2 sound is not "new," but rather "similar" or "identical" to the L1 sound. Given that transcription conventions for a given language are often based on phonemic considerations (e.g., the contrastive status of certain phonetic details), such a phonetic symbol test will often resemble a cross-linguistic phonemic analysis. Flege noted that the phonetic symbol test was not absolute, however, and that its results should be supplemented with acoustic and perceptual data in making predictions about the relation of L1 and L2 sounds. The shortcomings of this method were also pointed out by Bohn (2002), who noted that perceptual measures provide the most stable assessments of phonological similarity. Neither Flege nor Bohn specified how different types of similarity should be resolved when they make conflicting predictions.

The hypothesis examined here is that different types of cross-linguistic phonological similarity are resolved by L2 users in favor of higher-level similarity. In other words, the manner in which L2 users relate L2 segments to L1 segments is predicted to be based predominantly upon abstract, between-system comparisons at the phonemic level, not within-system comparisons at the allophonic level or system-external comparisons at a psychoacoustic level. As outlined above, such a cross-linguistic phonemic level of analysis probably considers at least a segment's position within relevant featural dimensions as well as distributional information. In this regard, it is important to point out that the term "L2 users" is meant to refer to individuals who would have access to this kind of information – that is, L2 users with phonemic knowledge of the L2, not naïve listeners exposed to the L2 for the first time. Relatively experienced L2 users are expected to show L1-L2 mappings that follow phonemic similarity over acoustic and allophonic similarity because of a tendency for high-level information to override low-level information,[5] consistent with many other "top-down" effects in speech processing (e.g., Warren 1970; Jongman *et al.* 2003; Davis and Johnsrude 2007). In the following section, these predictions are shown to be borne out in a wide range of cross-linguistic research.

3 Phonemic similarity in cross-linguistic research

3.1 Phonemic correspondence in second-language perception

Because languages differ in terms of phonemic inventory and patterns of allophonic alternation, both phonemic similarity and allophonic similarity are language-specific kinds of similarity, as explained in section 2.1. It should thus come as no surprise that L2 perception of a given phonological structure has been observed to vary across L1 backgrounds. The language-specific nature of cross-language mapping is often attributed to the existence of different phonological constraints, or different rankings of constraints, across languages (e.g., Broselow *et al.* 1998), but some part of this language specificity is likely due to basic cross-linguistic differences in the perception of unfamiliar phonological structures.

The Perceptual Assimilation Model (PAM; Best 1994, 1995) – an articulatory framework for understanding non-native speech perception – has played a particularly influential role in the analysis of cross-linguistic differences in non-native perception, attributing them to the various ways in which non-native sounds may align with the gestural constellations of native phonological categories. The core insight of the PAM lies in relating disparities in perception of foreign contrasts to disparities in native phonological knowledge gained through linguistic experience. Difficulty in discriminating a foreign contrast arises when the structure of a listener's native phonology interferes, causing the foreign sounds to be perceptually assimilated to a single native category. For example, when clicks – language sounds that are relatively distinct acoustically – are played to listeners with no native click categories, the clicks are discriminated well; in the case of click-language speakers, however, non-native clicks are discriminated poorly, due to convergent perceptual assimilation to native click categories (Best *et al.* 2003). Perceptual assimilation to L1 structures results in cross-linguistic differences not only in the discriminability of non-native segments, but also in the perception of non-native phonotactics, such as initial and medial consonant clusters (Dupoux *et al.* 1999; Hallé and Best 2007). These findings demonstrate that L1 phonological structure exerts a profound influence on cross-linguistic speech perception, biasing listeners of different language backgrounds toward interpreting the same L2 input in disparate ways (see also, e.g., Weinreich 1953; Flege 1987, 1995).

However, even though the specific percept of a given L2 segment may usually differ across L1 backgrounds, it is still reasonable to predict that the L2 segment will be perceptually assimilated to the L1 segment that is the closest *phonetically* (acoustically and/or articulatorily), whatever that may be. Consistent with this prediction, the literature on L2 speech perception includes many findings of close correspondence between acoustics, or phonetic realization more generally, and perceptual performance (see, e.g., Bohn and Best 2012 for recent findings on cross-linguistic discrimination of approximants). Nevertheless, some studies evince a dissociation between acoustic similarity and perceptual similarity. Perception of non-native vowels in particular has been repeatedly shown to bypass the acoustically closest L1 vowels in favor of the phonemically closest L1 vowels. For example, native speakers of Canadian English judge German /u/ to be more similar than German /y/ to English /u/, despite the fact that German /y/ is the acoustically closer vowel to English /u/ (Polka and Bohn 1996). Similarly, native speakers of American English judge front rounded vowels in both French and German to be more similar to English back rounded vowels, even though they are acoustically closer to the English front unrounded vowels in three-dimensional ($F_1 \times F_2 \times F_3$) vowel space (Strange *et al.* 2004). These findings demonstrate that cross-linguistic perceptual similarity does not follow straightforwardly from traditional measures of acoustic similarity; rather, listeners may perceive an L2 segment as most similar to the phonemically closest L1 segment, even when this is not the acoustically closest one.[6]

The pattern of perceiving the phonemically closest L1 segment as most similar to an L2 segment is consistent with the idea of perceptual assimilation at the phonological (abstract) level, as described in a version of the PAM for L2 learners, the PAM-L2 (Best and Tyler 2007). Although the PAM-L2 does not specify how the phonological level interacts with the phonetic level and the gestural level, it does state that this interaction is likely to change over time as L2 learners gain more knowledge of the L2, suggesting that phonemic information may come to play a larger role as more of it becomes available over the course of L2 learning. Indeed, when L2 learners have a modicum of L2 phonemic knowledge, they seem to prioritize phonemic correspondence over acoustic proximity in the calculation of overall cross-linguistic phonological similarity, as is evident in much of the work on L2 production.

3.2 Phonemic correspondence in second-language production

Like studies of L2 perception, studies of L2 production suggest that phonemic similarity plays a large role in relating L2 forms to L1 forms. In the Speech Learning Model (SLM; Flege 1995, 1996), a model of L2 perception and production that assumes an allophonic level of analysis, phonemic similarity is not discussed as such; however, this corresponds closely to what is measured in the phonetic symbol test (see section 2.3): phonemically similar sounds tend to be transcribed with the same symbol. Together with acoustic and perceptual similarity, phonemic similarity helps predict whether novel L2 sounds will be classified by L2 learners as "new," "similar," or "identical" with respect to familiar L1 sounds. Classification of an L2 sound as "identical" to an L1 category does not negatively impact its production, as any differences between the two are negligible, whereas classification of an L2 sound as "similar" to an L1 category negatively affects its production because of a significant phonetic disparity between the L1 and L2 sounds. Analogizing an L2 sound to a similar, but non-identical, L1 sound results in their perceptual linkage, which allows the disparate properties of the L1 sound to influence the production of the L2 sound (and vice versa). A "new" L2 sound, by contrast, has no clear L1 counterpart and is thus not analogized to an L1 sound, which allows it to be produced accurately, free from L1 influence, once sufficient experience in the L2 has been acquired.

Although the SLM, like the PAM(-L2), does not address the interaction among the various kinds of similarity that influence overall cross-linguistic similarity, the L2 speech literature implies that, for L2 learners, phonemic similarity takes precedence over acoustic phonetic similarity. Despite the tendency for different metrics of similarity to converge toward the same conclusions, it is not uncommon for phonemic similarity and acoustic similarity to be at odds with each other (see, e.g., Hammarberg 1996), as shown by the frequent disagreement between acoustic comparisons and transcription conventions, which for a given language are based partly on phonemic considerations. Since the SLM relates L1 and L2 sounds at a position-sensitive allophonic level, it seems to predict that L2 learners will resolve such conflicts in favor of acoustic phonetic comparisons; however, this is not how the "new" vs. "similar" distinction is applied throughout the literature, including Flege and Hillenbrand's (1984) study of American English speakers' production of French /u/ and French /y/

(in the words *tous* /tu/ "all" and *tu* /ty/ "you"). For L1 English speakers, L2 French /u/ is analyzed as a "similar" vowel with an L1 counterpart in English /u/, while L2 French /y/ is analyzed as a "new" vowel with no L1 counterpart. This classification contrasts with the acoustic facts, which show that in the given alveolar context, English /u/ is actually closer to French /y/ than to French /u/ (Strange *et al.* 2007). Nevertheless, the production of L2 French vowels by L1 English speakers is consistent with the phonemic pairing of vowels: French /y/ is produced accurately, while French /u/ shows influence from English /u/ – a result that has been replicated for L1 English learners of Mandarin (Chang *et al.* 2011).

Natural classes of consonants also show this bias toward phonemic correspondence in L2 production. For example, English-Portuguese bilinguals – influenced by the long-lag voice onset time (VOT) of L1 English voiceless stops – produce the L2 Portuguese voiceless unaspirated stops with relatively long VOT, not with the acoustically more similar short VOT characteristic of L1 English voiced stops (Major 1992). English-French bilinguals show the same effect with L2 French /t/, which they produce with VOT that is too long, under influence from long-lag L1 English /t/ (Flege 1987). In both cases, inauthentic, or accented, VOT production follows from a perceptual linkage of the L2 voiceless stops with the phonemically corresponding L1 voiceless stops, not with the L1 voiced stops (which, with their short-lag VOT, are arguably more similar at an acoustic phonetic level).[7] The pattern that emerges is that L2 users seem to favor linking L1 and L2 sounds on the basis of phonemic correspondence rather than strictly acoustic proximity.[8] This same pattern is found in cross-linguistic influence of the L2 on the L1.

3.3 Phonemic correspondence in phonetic drift

According to the SLM, the same mechanism of cross-linguistic perceptual linkage between "similar" L1 and L2 sounds is responsible for both L1-to-L2 interference and L2-to-L1 interference, although the theoretical distinction between "similar" and "new" L2 sounds has rarely been discussed in concrete acoustic terms (cf. Strange 1999 for a typology in terms of mappings and goodness-of-fit ratings). Chang (2010) suggested that, along with "identical" L2 sounds, "similar" and "new" L2 sounds form a continuum of cross-linguistic similarity to the L1, with the boundary between "identical" and "similar" generally corresponding to a cross-linguistic disparity of one just-noticeable difference (JND) in the relevant acoustic dimension. This conceptualization of the L2 in relation to the L1 is found to correlate well with patterns of L1 "phonetic drift" – subtle shifts in L1 speech production following from L2 experience in adulthood (e.g., Flege 1987; Major 1992; Sancier and Fowler 1997).

In a study of L1 English novice learners of Korean, Chang (2010, 2012a) showed that phonetic drift occurs after even brief L2 experience, at multiple levels of phonological structure, and in accordance with initial cross-linguistic distances between the L1 and the L2. Crucially, L1 structures drift toward L2 structures only when there is an appropriate amount of cross-linguistic distance between them. To trigger phonetic drift, an L2 structure must be similar enough to an L1 structure to qualify as "similar" rather than "new," yet different enough not to be perceived as "identical" (i.e., at least one JND away from the L1 in the relevant dimension). Thus, for L1 English learners of Korean, the VOT of L1 English voiceless stops drifted toward the VOT of the perceptually similar L2 Korean aspirated stops (which differs by more than the JND for VOT), whereas the VOT of L1 English voiced stops did not drift toward the VOT of the perceptually similar L2 Korean fortis stops (which differs by less than the JND for VOT).

Although the L1-L2 perceptual linkages proposed in Chang (2010, 2012a) are justified on an acoustic basis, these results do not provide clear evidence of a phonemic basis to relations between L1 and L2 sounds, since the relevant natural classes of English and Korean stops differ phonologically in a number of ways and the participants – novice learners – may not have had sufficient phonological knowledge of the L2 to draw phonemic comparisons in any case. Studies of phonetic drift in advanced L2 learners, however, are consistent in showing a preference for linking L1 and L2 sounds that correspond phonemically. For example, French-English bilinguals produce their short-lag L1 French /t/ with overly long VOT, under influence from long-lag L2 English /t/ (Flege 1987). From an acoustic perspective, it would be more favorable for them to link their L1 French /t/ instead to the similarly short-lag L2 English /d/, in which case they would not be expected to manifest much phonetic drift at all; however, the observed pattern of drift instead evinces a perceptual linkage between the phonemically corresponding voiceless stops. Similarly, Portuguese-English bilinguals produce L1 Portuguese voiceless stops with VOT that is influenced by the longer VOT of L2 English voiceless stops, not by the similar short-lag VOT of L2 English voiced stops (Sancier and Fowler 1997). These production data thus provide additional evidence of respect for phonemic correspondence between L1 and L2 stop types, a phenomenon that is also well documented in the literature on loanwords.

3.4 Phonemic correspondence in loanword phonology

The importance of phonemic similarity in determining cross-linguistic mappings has been amply demonstrated in studies on loanword adaptation, as reviewed by Kang (2008, 2011). This literature has shown that when foreign words are phonologically incorporated into a language, acoustic perceptual similarity interacts with phonemic similarity in complex ways. While many modifications to borrowed L2 forms mirror patterns of perceptual "foreign accent" (Peperkamp and Dupoux 2003), other changes do not seem to maximize perceptual similarity (e.g., Shinohara *et al.* 2011) and map an L2 form to the phonemically most similar L1 form instead of the perceptually most similar one. Similarity between L1 and L2 sounds at the phonemic level is able to play a role in loanword adaptation because the primary agents of adaptation are typically fluent bilinguals acquainted with the phonological structure of both the L1 and the L2, not monolingual L1 speakers (Paradis and LaCharité 1997, 2008). The outcome of loanword adaptation, however, is not fully determined by phonemics, as borrowed forms frequently evince a more fine-grained analysis of the L2 signal. Final voiced stops in English, for instance, are adapted into Korean variably as unreleased voiceless stops or lenis stops with a following epenthetic vowel, depending on the quality of the preceding vowel and the place of articulation of the stop (Kang 2003). This sort of variability reveals a nuanced sensitivity to the phonetics of the lending L2 that would be lost in a strictly phonemic analysis, suggesting that the phonemic representation of a borrowed word is enriched by phonetic detail (Kang 2008), which influences the outcome of adaptation at the same time as phonemic information (Chang 2012b).

Nevertheless, L2-to-L1 mapping in loanword adaptation often evinces a respect for source (L2) phonemics, and many relevant cases of following phonemic correspondence over acoustic proximity are reported in detail by LaCharité and Paradis (2005). Cross-linguistic mapping of vowels, for example, tends to occur on a phonemic basis in loanwords. In the case of English borrowings in Quebec French, the English high lax vowels /ɪ, ʊ/ are acoustically closest to the French mid vowels /ɛ, ɔ/, but these English vowels are consistently mapped to the French high vowels /i, u/, not to the French mid vowels /ɛ, ɔ/ (or /e, o/). In the case of English borrowings in Japanese, the English rhotic /r/ (realized as an alveolar approximant [ɹ]) is mapped onto the Japanese rhotic /r/ (realized as a postalveolar flap [ɾ̠]), not onto the Japanese approximant /w/, even though /w/ is perceptually more similar. Finally, in the case of English borrowings in Mexican Spanish, the English voiced stops, despite being realized as voiceless word-initially, are nearly always adapted as the strongly prevoiced Spanish voiced stops, not as the voiceless stops (which are more similar in terms of VOT). Similar findings are reported by Chang (2009a, 2012b) for English borrowings in Burmese. Even though English voiceless stops are typically realized as aspirated word-initially, they are nearly always borrowed into Burmese with the Burmese voiceless unaspirated stop series, not with the voiceless aspirated series. These data show overall a respect for source phonemic distinctions, which results in cross-linguistic mapping to phonemically parallel sounds even when these sounds are not the most similar acoustically.

4 Discussion

In this chapter, I have endeavored to make three points: (1) "phonological similarity" is a complex construct consisting of, and influenced by, multiple types of similarity; (2) levels of similarity are hierarchically organized, with high-level similarity ranked above low-level similarity; and (3) the influence of phonemic similarity, based on high-level information that is only relevant for cross-linguistic comparisons, is at least partly responsible for disparities between intra- and inter-language effects of low-level similarity. In cases of conflict, phonemic similarity tends to override acoustic perceptual similarity, with the result that cross-linguistic speech patterns often depart from predictions based on acoustic and/or perceptual similarities. As summarized in section 3, this trend is found in a range of cross-linguistic studies examining L2 perception, L2 production, L2-influenced phonetic drift in L1 production, and loanword adaptation.

Although I have argued that L2 users are swayed by phonemic correspondences when phonemic information is in conflict with low-level information, it is important to emphasize that this is a tendency, not a rule. In section 3.4, it was pointed out that loanword adaptation is not all about phonemic correspondences, and that the ultimate form of a loanword often bears traces of sensitivity to phonetic properties of the source language. For example, while Burmese adapts allophonically aspirated English voiceless stop allophones with unaspirated voiceless stops, Thai generally adapts these English allophones with aspirated voiceless stops rather than with unaspirated voiceless stops (Kenstowicz and Suchato 2006). The English-to-Thai mapping is thus an apparent counterexample to the ranking of phonemic similarity over acoustic phonetic similarity. However, when considered more carefully, the disparity between Burmese adaptation and Thai adaptation may actually be due to phonemic considerations after all. In Burmese, the adaptation of English voiceless stops with unaspirated voiceless stops allows aspirated voiceless stops (namely, /pʰ/) to be used to adapt certain English fricatives (namely, /f/) that are absent from the Burmese inventory, thus preventing phonemic contrasts between English fricatives and stops from being neutralized (Chang, 2012b). Thai also lacks certain English fricatives – namely, /θ/, which it adapts as /t/ (Kenstowicz and Suchato 2006). Adaptation of English voiceless stops with Thai unaspirated voiceless stops would, therefore, neutralize the contrast between English /θ/ and /t/, so instead they are adapted with Thai aspirated voiceless stops, which preserves the contrast between /θ/ and /t/.

In other words, although it is possible for "phonetic" kinds of similarity to prevail over "phonological" kinds of similarity in cases where they make different predictions, this appears to occur in extenuating circumstances having to do with other phonological considerations (or, alternatively, with an insufficient knowledge of the L2 phonology). As yet, it is not clear that an L2 user with phonemic knowledge of the L2 would ever weigh phonetic information *at the expense of* phonemic information (e.g., maximizing phonetic detail in a way that neutralizes phonemic contrast). The claim made here is that this is unlikely to happen because phonemic similarity has a privileged status stemming from its connection to establishing and maintaining meaningful linguistic contrast.

In section 3.1 it was observed that cross-linguistic similarity between segments differs from within-language similarity between segments in two ways: the relevance of between-system comparisons at a phonemic level, which are applicable only in cross-linguistic situations, and the language-specific nature of cross-linguistic perceptual similarity, which arises due to cross-linguistic differences in the landscape of L1 "perceptual magnets" (Kuhl and Iverson 1995) for unfamiliar L2 sounds. When L2 phonemic information is available, it exerts a powerful influence on cross-linguistic segmental mapping that can override conflicting information from acoustic phonetic similarity. In this sense, phonemic similarity constitutes one of multiple factors that may mask effects of "raw" perceptual similarity between languages. As discussed by Shaw and Davidson (2011), recoverability and uniformity are other factors that may interact with perceptual similarity in determining the output of cross-linguistic production. The challenge for future cross-linguistic speech research will be to account for how much cross-linguistic differences in the grammatical effects of perceptual similarity have to do with variation in the "P-maps" of speakers of diverse languages (due to the perceptual warping caused by linguistic knowledge) vs. other impinging factors, such as abstract phonemic comparisons.

5 Conclusion

The research findings reviewed in this chapter suggest that the way in which L2 segments are related to L1 segments differs fundamentally from the way in which L1 segments are related to other L1 segments. I have argued that this disparity is rooted in a phonemic level of segmental comparisons that is only relevant between languages. Phonemic similarity distinguishes cross-linguistic phonological similarity from within-language phonological similarity because only judgments of cross-linguistic similarity can be influenced by between-system analyses of two phonologies. To the extent that such phonemic comparisons may depart from acoustic and allophonic comparisons, the availability of this high-level information can lead to the appearance that low-level information is being ignored, since high-level information is likely to prevail in cases of conflict.

Although the studies discussed in this chapter provide evidence for the privileged status of phonemics in determining overall phonological similarity between L1 and L2 segments, it is logical to expect differences between L2 learners, who have phonemic knowledge of the L2, and naïve non-natives, who do not. If we can assume, as implied by Flege (1995), that L2 sounds undergo automatic equivalence classification with L1 sounds, this suggests that at the onset of L2 learning L1 and L2 sounds must be linked on the basis of low-level information. It is clear that perceived cross-linguistic similarity based on this kind of low-level information is related to cross-linguistic behavior (e.g., Baker *et al.*, 2002; Aoyama *et al.*, 2004), but what remains unclear is how perceived similarity between L1 and L2 sounds changes over the course of L2 learning. This question motivates several interesting avenues of future research into the effects of L2 phonemic information over time and the manner in which a changing level of cross-linguistic linkage modulates L1-to-L2 and L2-to-L1 influence as an L1 talker acquires an L2 phonology.

Acknowledgments

I am thankful to the audience at the CUNY Phonology Forum (especially William Idsardi and Douglas Whalen) and to the editors of this volume, Charles Cairns and Eric Raimy, for their insightful feedback on this paper. Any errors are mine alone.

Notes

1 Of course, the term "hierarchical" can refer to a variety of systems, ranging from a strictly feed-forward system to one in which information flows freely between modules. While the precise nature of the relationship between phonetics and phonology lies outside the scope of this chapter, the crucial point is that phonology is privileged (higher-ranked) relative to phonetics.

2 I say "at least" because I have clearly omitted other dimensions of phonological similarity (e.g., articulatory/gestural similarity, aerodynamic similarity) in the interest of focusing on the contrast between low-level and high-level similarity.

3 I am conflating acoustic and auditory similarity here since the distinction is not important for the main contrast between "phonetic" and "phonological" kinds of similarity. However, it is important to note that acoustic and auditory similarity are in fact different. While auditory impression may broadly reflect the acoustics of the speech signal, auditory distances are not linearly related to acoustic distances because of basic aspects of hearing and auditory processing such as the pattern of frequency response of the basilar membrane in the inner ear. This fact does not affect the main argument; nevertheless, it should be borne in mind that although the phonetic distances referred to throughout this chapter are acoustic, the relevant distances in regard to perceptual similarity are really auditory.

4 The reason for supplementing phonological features with the notion of "relative phonetics" (by which I mean "relative position in a relevant phonetic dimension"; e.g., "long" end of the voice onset time dimension) is to ensure a common currency of comparison between L1 and L2 sounds, which feature specifications may not always provide. For example, stop laryngeal categories – whether "voicing" or "aspiration" categories – have been widely described phonetically in terms of the acoustic property of voice onset time, but phonologically in terms of at least two different features, [±voice] and [±spread glottis], depending on whether the contrast is one of voicing or aspiration. Relating stop types in a "voicing" language to those in an "aspirating" language via feature matching is, therefore, problematic; however, doing so in terms of relative phonetics is straightforward, since relative position in the voice onset time dimension is something that can be meaningfully compared for stop types in different kinds of languages.

5 Another prediction that follows from this hypothesis is that in case of a conflict between acoustic similarity and allophonic similarity, listeners will, depending on the nature of the task, be swayed by allophonic similarity over acoustic similarity in determining overall perceptual similarity between a pair of segments. That is to say, listeners whose native language contains a productive alternation only between a pair of phones that are relatively acoustically distinct (e.g., [s] and [θ]) are expected to perceive that pair of phones as more similar than a pair of phones that are acoustically closer (e.g., [f] and [θ]) but do not participate in such an alternation. This seems to be a reasonable prediction, but for reasons of space the discussion below is limited to conflicts between acoustic similarity and phonemic similarity.

6 Acoustic proximity in these studies has generally been measured in terms of distance in the first two or three vowel formants (F_1–F_3). However, there are limits to estimating acoustic proximity in these terms, since F_1–F_3, though sufficient as acoustic cues for distinguishing most vowels, are not the only determinants of vowel quality. Thus, it should be noted that inclusion of additional acoustic dimensions – especially the fundamental frequency (f_0) and the temporal trajectories of frequency components – would give a

fuller picture of acoustic proximity between vowels and may help account for perceptual assimilations to an L1 vowel that is not the closest to an L2 vowel as measured on the basis of F_1–F_3 alone.

7 Like all phonological contrasts, voicing contrast is associated with multiple acoustic cues besides VOT (e.g., f_0 in the adjacent vowel). However, there is evidence that English speakers rely on VOT as the primary cue to voicing (Abramson and Lisker 1985), suggesting that for English speakers perceptual similarity between L1 and L2 stop types is likely to closely follow their VOT characteristics.

8 Naturally, many of these cross-linguistic linkages between segments argued to be based on phonemic similarity could also be explained in terms of orthographic similarity (i.e., being spelled with the same graphemes). However, even when the L1 and L2 share the same alphabet, not all of the data can be explained in this way; see LaCharité and Paradis (2005, pp. 240–241) for extensive arguments against attributing cross-linguistic mappings to orthographic influence.

References

Abramson, A. and L. Lisker (1985). Relative power of cues: F0 shift versus voice timing. In V. Fromkin (ed.), *Phonetic Linguistics: Essays in Honor of Peter Ladefoged*. San Diego, CA: Academic Press, pp. 25–33.

Aoyama, K., J. E. Flege, S.G. Guion, R. Akahane-Yamada, and T. Yamada (2004). Perceived phonetic dissimilarity and L2 speech learning: the case of Japanese /r/ and English /l/ and /r/. *Journal of Phonetics* 32(2): 233–250.

Austin, W.M. (1957). Criteria for phonetic similarity. *Language* 33(4): 538–544.

Baker, W., P. Trofimovich, M. Mack, and J.E. Flege (2002). The effect of perceived phonetic similarity on non-native sound learning by children and adults. In B. Skarabela, S. Fish, and A.H.-J. Do (eds.), *Proceedings of the Twenty-sixth Annual Boston University Conference on Language Development*, vol. 1. Somerville, MA: Cascadilla Press, pp. 36–47.

Best, C.T. (1994). The emergence of native-language phonological influences in infants: a perceptual assimilation model. In J.C. Goodman and H.C. Nusbaum (eds.), *The Development of Speech Perception: The Transition from Speech Sounds to Spoken Words*. Cambridge, MA: MIT Press, pp. 167–224.

Best, C.T. (1995). A direct realist view of cross-language speech perception. In W. Strange (ed.), *Speech Perception and Linguistic Experience: Issues in Cross-Language Research*. Baltimore, MD: York Press, pp. 171–204.

Best, C.T., A. Traill, A. Carter, K.D. Harrison, and A. Faber (2003). !Xóõ click perception by English, Isizulu, and Sesotho listeners. In M.-J. Solé, D. Recasens, and J. Romero (eds.), *Proceedings of the Fifteenth International Congress of Phonetic Sciences*. Barcelona, Spain: Causal Productions, pp. 853–856.

Best, C.T. and M.D. Tyler (2007). Nonnative and second-language speech perception: commonalities and complementarities. In O.-S. Bohn and M.J. Munro (eds.), *Language Experience in Second Language Speech Learning: In Honor of James Emil Flege*. Amsterdam: Benjamins.

Bohn, O.-S. (2002). On phonetic similarity. In P. Burmeister, T. Piske, and A. Rohde (eds.), *An Integrated View of Language Development: Papers in Honor of Henning Wode*. Trier: Wissenschaftlicher Verlag, pp. 191–216.

Bohn, O.-S. and C.T. Best (2012). Native-language phonetic and phonological influences on perception of American English approximants by Danish and German listeners. *Journal of Phonetics* 40(1): 109–128.

Boomershine, A., K.C. Hall, E. Hume, and K. Johnson (2008). The impact of allophony versus contrast on speech perception. In P. Avery, E. Dresher, and K. Rice (eds.), *Contrast in Phonology*. Berlin: Mouton, pp. 143–172.

Broselow, E., S.-I. Chen, and C. Wang (1998). The emergence of the unmarked in second language phonology. *Studies in Second Language Acquisition* 20(2): 261–280.

Chang, C.B. (2009a). English loanword adaptation in Burmese. *Journal of the Southeast Asian Linguistics Society* 1: 77–94.

Chang, C.B. (2009b). The status of voicing and aspiration as cues to Korean laryngeal contrast. In M. Elliott, J. Kirby, O. Sawada, E. Staraki, and S. Yoon (eds.), *Proceedings from the Forty-third Annual Meeting of the Chicago Linguistic Society: The Main Session*. Chicago, IL: Chicago Linguistic Society, pp. 31–45.

Chang, C.B. (2010). First language phonetic drift during second language acquisition. Ph.D. dissertation, University of California Berkeley.

Chang, C.B. (2012a). Rapid and multifaceted effects of second-language learning on first-language speech production. *Journal of Phonetics* 40(2): 249–268.

Chang, C.B. (2012b). Phonetics vs. phonology in loanword adaptation: revisiting the role of the bilingual. In S. Berson, A. Bratkievich, D. Bruhn, A. Campbell, R. Escamilla, A. Giovine, L. Newbold, M. Perez, M. Piqueras-Brunet, and R. Rhomieux (eds.), *Proceedings of the Thirty-fourth Annual Meeting of the Berkeley Linguistics Society: General Session and Parasession on Information Structure*. Berkeley, CA: Berkeley Linguistics Society, pp. 61–72.

Chang, C.B., Y. Yao, E.F. Haynes, and R. Rhodes (2011). Production of phonetic and phonological contrast by heritage speakers of Mandarin. *Journal of the Acoustical Society of America* 129(6): 3964–3980.

Copeland, D.E. and G.A. Radvansky (2001). Phonological similarity in working memory. *Memory and Cognition* 29(5): 774–776.

Covington, M.A. (1996). An algorithm to align words for historical comparison. *Computational Linguistics* 22(4): 481–496.

Davis, M.H. and I.S. Johnsrude (2007). Hearing speech sounds: top-down influences on the interface between audition and speech perception. *Hearing Research* 229(1–2): 132–147.

Dupoux, E., Y. Hirose, K. Kakehi, C. Pallier, and J. Mehler (1999). Epenthetic vowels in Japanese: a perceptual illusion? *Journal of Experimental Psychology: Human Perception and Performance* 25(6): 1568–1578.

Flege, J.E. (1987). The production of "new" and "similar" phones in a foreign language: Evidence for the effect of equivalence classification. *Journal of Phonetics* 15(1): 47–65.

Flege, J.E. (1995). Second language speech learning: theory, findings, and problems. In W. Strange (ed.), *Speech Perception and Linguistic Experience: Issues in Cross-Language Research*. Baltimore, MD: York Press, pp. 233–272.

Flege, J.E. (1996). English vowel productions by Dutch talkers: more evidence for the "similar" vs "new" distinction. In A. James and J. Leather (eds.), *Second-Language Speech: Structure and Process*. Berlin, Germany: Mouton, pp. 11–52.

Flege, J.E. and J. Hillenbrand (1984). Limits on phonetic accuracy in foreign language speech production. *Journal of the Acoustical Society of America* 76(3): 708–721.

Flemming, E. (2004). Contrast and perceptual distinctiveness. In B. Hayes, R. Kirchner and D. Steriade (eds.), *Phonetically-Based Phonology*. Cambridge: Cambridge University Press, pp. 232–276.

Fournet, N., A. Juphard, C. Monnier, and J.-L. Roulin (2003). Phonological similarity in free and serial recall: the effect of increasing retention intervals. *International Journal of Psychology* 38(6): 384–389.

Frisch, S.A., J.B. Pierrehumbert, and M.B. Broe (2004). Similarity avoidance and the OCP. *Natural Language and Linguistic Theory* 22(1): 179–228.

Gahl, S., Y. Yao, and K. Johnson (2012). Why reduce? Phonological neighborhood density and phonetic reduction in spontaneous speech. *Journal of Memory and Language* 66(4): 789–806.

Gallagher, G. (2012). Perceptual similarity in non-local laryngeal restrictions. *Lingua* 122(2): 112–124.

Gallagher, G. and P. Graff (2012). The role of similarity in phonology. *Lingua* 122(2): 107–111.

Hallé, P.A. and C.T. Best (2007). Dental-to-velar perceptual assimilation: a cross-linguistic study of the perception of dental stop+/l/ clusters. *Journal of the Acoustical Society of America* 121(5): 2899–2914.

Hammarberg, B. (1996). Conditions on transfer in phonology. In A. James and J. Leather (eds.), *Second-Language Speech: Structure and Process*. Berlin: Mouton, pp. 161–180.

Johnson, K. (2003). *Acoustic and Auditory Phonetics* (2nd edition). Oxford: Blackwell.

Johnson, K. and M. Babel (2010). On the perceptual basis of distinctive features: Evidence from the perception of fricatives by Dutch and English speakers. *Journal of Phonetics* 38(1): 127–136.

Jongman, A., Y. Wang, and B.H. Kim (2003). Contributions of semantic and facial information to perception of nonsibilant fricatives. *Journal of Speech, Language, and Hearing Research* 46(6): 1367–1377.

Kang, Y. (2003). Perceptual similarity in loanword adaptation: English postvocalic word-final stops in Korean. *Phonology* 20(2): 219–273.

Kang, Y. (2008). Interlanguage segmental mapping as evidence for the nature of lexical representation. *Language and Linguistics Compass* 2(1): 103–118.

Kang, Y. (2011). Loanword phonology. In M. van Oostendorp, C.J. Ewen, E. Hume, and K. Rice (eds.), *The Blackwell Companion to Phonology*, vol. IV: Phonological Interfaces. Oxford: Wiley-Blackwell, pp. 2258–2282.

Kawahara, S. and K. Garvey (2010). Testing the P-map hypothesis: coda devoicing. *Rutgers Optimality Archive* 1087.

Kenstowicz, M. and A. Suchato (2006). Issues in loanword adaptation: a case study from Thai. *Lingua* 116(7): 921–949.

Kessler, B. (2005). Phonetic comparison algorithms. *Transactions of the Philological Society* 103(2): 243–260.

Kondrak, G. (2003). Phonetic alignment and similarity. *Computers and the Humanities* 37(3): 273–291.

Kuhl, P.K. and P. Iverson (1995). Linguistic experience and the "perceptual magnet effect." In W. Strange (ed.), *Speech Perception and Linguistic Experience: Issues in Cross-Language Research*. Baltimore, MD: York Press, pp. 121–154.

LaCharité, D. and C. Paradis (2005). Category preservation and proximity versus phonetic approximation in loanword adaptation. *Linguistic Inquiry* 36(2): 223–258.

Ladefoged, P. (1969). The measurement of phonetic similarity. In *COLING '69: Proceedings of the 1969 Conference on Computational Linguistics*. Stroudsburg, PA: Association for Computational Linguistics, pp. 1–14.

Laeufer, C. (1996). Towards a typology of bilingual phonological systems. In A. James and J. Leather (eds.), *Second-Language Speech: Structure and Process*. Berlin: Mouton, pp. 325–342.

Luce, P.A. and D.B. Pisoni (1998). Recognizing spoken words: the neighborhood activation model. *Ear and Hearing* 19(1): 1–36.

Major, R.C. (1987). Phonological similarity, markedness, and rate of L2 acquisition. *Studies in Second Language Acquisition* 9(1): 63–82.

Major, R.C. (1992). Losing English as a first language. *The Modern Language Journal* 76(2): 190–208.

Mielke, J. (2012). A phonetically based metric of sound similarity. *Lingua* 122(2): 145–163.

Miller, G.A. and P.E. Nicely (1955). An analysis of perceptual confusions among some English consonants. *Journal of the Acoustical Society of America* 27(2): 338–352.

Munson, B. and N.P. Solomon (2004). The effect of phonological neighborhood density on vowel articulation. *Journal of Speech, Language, and Hearing Research* 47(5): 1048–1058.

Page, M.P.A., A. Madge, N. Cumming, and D.G. Norris (2007). Speech errors and the phonological similarity effect in short-term memory: evidence suggesting a common locus. *Journal of Memory and Language* 56(1): 49–64.

Paradis, C. and D. LaCharité (1997). Preservation and minimality in loanword adaptation. *Journal of Linguistics* 33(2): 379–430.

Paradis, C. and D. LaCharité (2008). Apparent phonetic approximation: English loanwords in Old Quebec French. *Journal of Linguistics* 44(1): 87–128.

Peperkamp, S. and E. Dupoux (2003). Reinterpreting loanword adaptations: the role of perception. In M.-J. Solé, D. Recasens, and J. Romero (eds.), *Proceedings of the Fifteenth International Congress of Phonetic Sciences*. Barcelona: Causal Productions, pp. 367–370.

Pierrehumbert, J. (1993). Dissimilarity in the Arabic verbal roots. In A. Schafer (ed.), *Proceedings of the Twenty-third Meeting of the North East Linguistic Society*. Amherst, MA: Graduate Linguistics Students Association, pp. 367–381.

Polka, L. and O.-S. Bohn (1996). A cross-language comparison of vowel perception in English-learning and German-learning infants. *Journal of the Acoustical Society of America* 100(1): 577–592.

Sancier, M.L. and C.A. Fowler (1997). Gestural drift in a bilingual speaker of Brazilian Portuguese and English. *Journal of Phonetics* 27(4): 421–436.

Shaw, J.A. and L. Davidson (2011). Perceptual similarity in input–output mappings: a computational/experimental study of non-native speech production. *Lingua* 121(8): 1344–1358.

Shinohara, S., S.-R. Ji, T. Ooigawa, and T. Shinya (2011). The limited role of perception in Korean loanword adaptation: the Korean three-way laryngeal categorization of Japanese, French, English and Chinese plosives. *Lingua* 121(9): 1461–1484.

Sloutsky, V.M. and A.V. Fisher (2012). Linguistic labels: conceptual markers or object features? *Journal of Experimental Child Psychology* 111(1): 65–86.

Steriade, D. (2009). The phonology of perceptibility effects: the P-map and its consequences for constraint organization. In K. Hanson and S. Inkelas (eds.), *The Nature of the Word: Studies in Honor of Paul Kiparsky*. Cambridge, MA: MIT Press, pp. 151–179.

Strange, W. (1999). Levels of abstraction in characterizing cross-language phonetic similarity. In J.J. Ohala, Y. Hasegawa, M. Ohala, D. Granville, and A.C. Bailey (eds.), *Proceedings of the Fourteenth International Congress of Phonetic Sciences*. New York: American Institute of Physics, pp. 2513–2519.

Strange, W., E. Levy, and R. Lehnholf, Jr. (2004). Perceptual assimilation of French and German vowels by American English monolinguals: acoustic similarity does not predict perceptual similarity. *Journal of the Acoustical Society of America* 115(5): 2606.

Strange, W., A. Weber, E.S. Levy, V. Shafiro, M. Hisagi, and K. Nishi (2007). Acoustic variability within and across German, French and American English vowels: phonetic context effects. *Journal of the Acoustical Society of America* 122(2): 1111–1129.

Vitevitch, M.S. (2002). The influence of phonological similarity neighborhoods on speech production. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 28(4): 735–747.

Warren, R.M. (1970). Restoration of missing speech sounds. *Science* 167(3917): 392– 393.

Weinreich, U. (1953). *Languages in Contact: Findings and Problems*. New York: Linguistic Circle of New York.

Zwicky, A.M. and E.D. Zwicky (1986). Imperfect puns, markedness, and phonological similarity: with fronds like these, who needs anemones? *Folia Linguistica* 20(3/4): 493–503.