Perceptual repair and syllable structure: a reply to Kabak and Idsardi

Donca Steriade, MIT

Address for correspondence:

Donca Steriade
Linguistics and Philosophy
32D828
Massachusetts Institute of Technology
77 Massachusetts Ave
Cambridge, MA 02139

Phone: 1-617-253-3129
Email: steriade@mit.edu

Acknowledgments

Many thanks to Barış Kabak for sharing with me a pre-publication version of Kabak and Idsardi (in press), and to Hyesun Cho, Jongho Jun, Yoonjung Kang and Heesun Kim for advice and insights on Korean phonology.

1. Introduction

Kabak and Idsardi (in press) report the results of experiments on the Korean speakers' perception of non-native CC clusters. One result is that Korean speakers cannot distinguish C_1C_2 - C_1VC_2 pairs, whose C_1 is an affricate like [tʃ]. This looks like an instance of the perceptual illusion of epenthesis reported earlier for Japanese by Dupoux et al. 1999. The Korean subjects *can* distinguish other C_1C_2 - C_1VC_2 pairs, where C_1C_2 is non-native. The authors base on these results an argument for the distinction between syllable-sensitive and syllable-insensitive phonotactics: only former are able trigger perceptual epenthesis. This note outlines an alternative interpretation, which is compatible with a wider range of evidence from Korean L2 phonology but lacks the broader implications of Kabak and Idsardi's own analysis, and offers a different perspective on other general issues brought up by the authors.

2. The experiments

Kabak and Idsardi report that Korean subjects discriminate the following pairs of sequences, all presented in V_V contexts, about as well as English speakers can:

(1) C_1C_2 whose C_1 is a legal Korean coda and C_2 a legal onset[1]

km	kʊm
gm	gʊm
gt^h	$gʊt^h$
ln	nn, ll

The C_1C_2 sequences in (1) are illegal but can be parsed into well formed syllables. Those in (2), below, are both illegal and contain ill-formed codas: only stops, nasals and [l] can end a Korean syllable.

[1] Voiced stops like [g] are identified as possible codas when the redundant feature of voicing is ignored.

(2) C_1C_2 whose C_1 is an illegal Korean coda

tʃt^h	tʃit^h
dʒt^h	dʒit^h
tʃm	tʃim
dʒm	dʒim

Korean subjects discriminated the pairs in (2) at substantially inferior levels compared to English speakers and compared to their performance on (1).

3. Potential significance

The comparison of (1) and (2) could bear on a number of questions. First, it could represent an extension of Dupoux et al.'s (1999, 2001)'s findings about illusory vowel effects to the different phonotactics of Korean. The result for Japanese is that none of the heterorganic C_1C_2 clusters tested, including versions of the clusters in (1), were discriminated from their C_1VC_2 modifications. What is the difference then between Japanese and Korean perception of non-native clusters? Kabak and Idsardi propose that all heterorganic C_1C_2 clusters in Japanese are syllabically ill formed, just like the Korean clusters in (2), because they all contain impossible codas: the only legal Japanese coda is the first half of a geminate or a nasal homorganic to a following obstruent (Ito 1988). Then the perception by Japanese subjects of heterorganic C_1C_2 should be compared only to that of Korean sequences in (2), which also contain illegal codas.

On the same interpretation, the Korean experiments identify a difference *within* the phonotactic constraints of Korean: the syllabic constraints – e.g. *strident-in-Coda - are able to skew perception, but constraints of the form *C_iC_j – which are independent of syllable position - are not. This distinction cannot be made in Japanese.

Finally, the experiments might tell us something about feature mismatches between the stimulus and a corresponding lexical entry: clusters like [gm], [gt^h] contained voiced codas which

impossible in Korean, but they were better discriminated from their C_1VC_2 version than clusters like [tʃt]. This finding is interpreted through the assumption that redundant features (here, voicing) can be ignored in perception, while distinctive ones (here, stridency) are not.

4. Confounds and an alternative

I outline now an interpretation of Kabak and Idsardi's results under which the data does not bear on the issues they highlight.

4.1. Illusory vowels and native epenthesis

In Dupoux, Kakehi, Hirose, Pallier and Mehler's (1999) experiments, Japanese subjects did not discriminate C_1ɯC_2 from C_1C_2. The Japanese vowel [ɯ] is the default epenthetic vowel of the language[2]. A later study by Dupoux, Pallier, Kakehi and Mehler (2001), using a lexical decision task, found differences between C_1ɯC_2 - C_1C_2 and C_1aC_2 - C_1C_2: stimuli containing C_1C_2 that had a lexical C_1ɯC_2 counterpart were judged to be real words on a majority of trials, as though [ɯ] had been perceptually inserted in C_1C_2; while C_1C_2-stimuli whose lexical counterpart contained C_1aC_2 (e.g. *mikdo* – related to real *mikado*) were rejected. The contrast between C_1ɯC_2 - C_1C_2 and C_1aC_2 - C_1C_2 suggests a link between the epenthetic status of [ɯ] in the phonology of Japanese and the Japanese listeners' propensity to perceive it in CC clusters. A version of this link was formulated by Dupoux et al.(1999) and Pepperkamp and Dupoux (2003): C_1ɯC_2 - C_1C_2 pairs are perceptually closer than C_1aC_2 - C_1C_2, perhaps because [ɯ] is shorter than [a], and this same factor explains why [ɯ] is the epenthetic vowel in the first place[3].

---

[2] Ito and Mester 1995. Note that Japanese epenthesizes [i], not [ɯ], after [ʃ], [tʃ]. The stimuli of Dupoux et al's (1998) experiments 2-4 contain 4 stimuli with ʃC clusters, from a total of 16 C_1C_2 stimuli. There was no comparison provided between the perception of ʃC clusters and that of others. In the later study by Dehaene-Lambertz, Dupoux and Gout (2000), the ʃC clusters were eliminated, with the suggestion (fn. 3) that the results they had provided in the original 1998 experiments had been less robust that those for other C_1C_2. Our conjecture is that the ʃɯC-ʃC pairs – like all pairs differentiated by a non-epenthetic vowel – would be more discriminable than C_1VC_2- C_1C_2 pairs where V is epenthetic.

[3] Cf. Steriade (in press) on the link between the typology of epenthetic choice and the phonetic properties of epenthetic vowels.

Dupoux and Pepperkamp (2002) also report that Brasilian Portuguese speakers – whose language epenthesizes [i] in illegal clusters – had error rates in the discrimination of C_1iC_2 - C_1C_2 comparable to those of $C_1ɯC_2$ - C_1C_2 for Japanese, lending further support to identification of the illusory vowel as the epenthetic vowel of the language.

Having suggested that pairs like C_1VC_2-C_1C_2 are hard to discriminate insofar as V is epenthetic, we turn to the results of Kabak and Idsardi's Korean experiments, which also employed pairs of the form C_1VC_2-C_1C_2. V had the appropriate epenthetic quality, [i], for the clusters in (2), where C_1= [tʃ], [dʒ]. These consonants trigger epenthesis of [i] in Korean loans.

(3) [i]-epenthesis after affricates

a. orange [orendʒi]
b. lunch [rɔntʃi]

In the case of the clusters in (1), however, Korean epenthesizes [ɨ] (Mira Oh p.c.)[4]:

(4) [ɨ]-epenthesis after stops

a. cockney [k^{h}ok^{h}ɨni]
b. Vietnam [pet^{h}ɨnam]
c. acme [ek^{h}mi]
d. litmus [rit^{h}ɨmosɨ]
e. Katmandu [k^{h}ɛt^{h}ɨmandu]

Now, for the stop-C clusters in (1), Kabak and Idsardi's subjects had to discriminate the pairs C_1C_2-C_1[ʊ]C_2, and not C_1C_2-C_1[ɨ]C_2: for instance, they had to compare [p^{h}agma] to [p^{h}agʊma], not [p^{h}agɨma]. English [ʊ] overlaps Korean [ɨ] on some acoustic dimensions, as

[4] A subset of these clusters – [km], [kn] - can also be nativized in Korean by regressive nasalization, in more frequent words. This does not extend to the t-nasal sequences and, in any case, what is relevant here is that any epenthetic vowel used in this case would be [ɨ].

Kabak and Idsardi note, but it appears to be perceived by Korean speakers as a variant of [u], judging from the fact it is invariably borrowed as [u], not [ɨ] (Hye-Sun Cho, Jongho Jun, Yoonjung Kang, Hee-Sun Kim p.c.)[5]:

(5) English [ʊ] is nativized as Korean [u]

f. full [p^{h}ul]

g. hook [huk, huk^{h}ɨ]

h. foot [p^{h}ut]

i. look [luk]

Judging from this data, the Korean subjects who heard [p^{h}agʊma] were likely to have represented it as [p^{h}aguma].

Now recall the results reported: Korean subjects successfully discriminated the C_1C_2-C_1[ʊ]C_2pairs in (1) but not the C_1C_2-C_1[i]C_2 pairs in (2), where C_1 is [tʃ], [dʒ]. A simple interpretation of this finding is that the subjects perceive as identical those C_1C_2-C_1VC_2 pairs that are related by Korean epenthesis – so [p^{h}atʃma]-[p^{h}atʃima] – but not the C_1C_2-C_1VC_2 pairs where V is non-epenthetic: that's exactly what happened, mutatis mutandis, in Japanese (Dupoux et al. 2001).

The use of C_1C_2-C_1[ʊ]C_2 pairs makes it hard to interpret the comparison between Korean and Japanese: we don't know if the Korean subjects were able to distinguish C_1C_2-C_1[ʊ]C_2 because the constraints violated by C_1C_2 are non-syllabic, unlike in Japanese, or because [ʊ] is not epenthetic in Korean, while the [ɯ] of the Japanese pairs C_1C_2-C_1[ɯ]C_2 is epenthetic in that language.

[5] I am informed by Hye-Sun Cho that further details on the perception of [ʊ] as [u] are found in Jun 2005. The borrowing of [ʊ] as [u] is not a spelling effect: English [ʌ], spelled <u>, as in *bus, cut*, is borrowed as [ɔ].

4.2. Phonotactics

The subjects in Kabak and Idsardi's study were Korean students residing in the US and thus likely to be fluent in English. In the instructions that preceded the experiment, they were told that "they would hear an American man saying nonsense words of English in pairs" (p. 23). Thus the phonotactic expectations these subjects possess *for English* may be as relevant to the outcome of the experiment as their Korean linguistic background.

A phonotactic constraint of English is that affricate-C sequences are prohibited morpheme internally. Affricate-C clusters occur across boundaries (e.g. *judg-ment, rich-ly*) but the instructions to the experiment and the stimuli, as they unfolded, provided no evidence of morphemic complexity. Therefore, when told that they will hear "nonsense words of English", the Korean subjects may have been indirectly led to expect sequences satisfying the morpheme-internal phonotactics of English. These allow stop-C, but not affricate-C sequences. The clusters in (2) - [tʃm], [dʒm], [tʃt^h], [dʒt^h]- contradicted these expectations: the words displayed affricate-C clusters with no evidence of a morphemic boundary separating them. By contrast, the clusters in (1) are well-formed morpheme medially, in English: cf. *picknick, enigma, pragmatic, pygmy*.

4.3. The alternative

We have seen that the sequences in (2), e.g. [padʒma]-[padʒima], which Korean subjects found hard to discriminate, cumulate three properties:

(a) the CC clusters contain impossible codas in Korean.
(b) the clusters are impossible inside English morphemes,
(c) their CVC transforms were separated by a vowel that's epenthetic in Korean

The easy-to-discriminate sequences in (1) differ from them on all three counts:

(d) the CC clusters contain legal codas in Korean.
(e) the clusters are acceptable in and across English morphemes,
(f) their CVC transforms were separated by a vowel that's not epenthetic in Korean

Kabak and Idsardi focus on factor (a), but to reach their conclusions it must be shown that (b) and (c) are irrelevant. In fact both (b) and (c) are very likely relevant, on the following scenario. The subjects heard words spoken by an American. The words could in principle have been represented as Korean or as English nonce words. They contained sometimes violations of Korean phonotactics (e.g. *pakma*) and sometimes violations of both English and Korean phonotactics (e.g. p*atʃma*). The subjects could react in several ways to a perceived violation of Korean phonotactics: they could decide that a word that is ill formed in Korean is well formed in English and they could assign it a representation on that basis; failing that, they could "correct" the raw percept, via epenthesis or in other ways. The assumption that they heard an English word would help with the sequences like *pakma*: the subjects could assign it a representation that is well-formed in English and faithfully store it as the English stimulus they had heard. Then they would be able to correctly discriminate it from any epenthetic transform presented later. On the other hand, sequences like *patʃma* could not be represented as English words. Unable to provide a phonological representation for the word in either language, the subjects could resort to a form of perceptual repair, here epenthesis, and store the stimulus *patʃma* as *patʃima*. This would make them unable to discriminate *patʃma*, stored as *patʃima*, from its epenthetic transform *patʃima*.

On this scenario, it matters whether the vowel is interpretable as epenthetic or not, because the subject who resorts to perceptual epenthesis is in effect mentally applying the native epenthesis process as he commits forms like *patʃma* to memory. Had the speaker been asked to discriminate *patʃma* not from *patʃima* but from *patʃʊma*, the results would have been different: *patʃma* would be stored as *patʃ**i**ma*, which would then be successfully discriminated

from *patʃʊma*. So it matters that in the pairs in (2) - *patʃma* vs *patʃima* - [i] is epenthetic: a different vowel would not have yielded the same effect.

Now suppose that the subject is asked to discriminate, as in the current experiment, *pakma* from *pakʊma*. One possibility is that he stores the form faithfully as *pakma,* recognizing [km] as a legal sequence for English. Then the stimulus *pakʊma* will be discriminated from the stored form *pakma*. Suppose, alternatively, that the Korean subject is unable to provide a phonotactically legal structure for *pakma*, perhaps because of uncertainty about the status of stop-nasal clusters in English. He will then apply perceptual epenthesis and store *pakma* as *pakɨma.* Now he has to discriminate this first stimulus from *pakʊma*: this is again possible, because he remembers the first stimulus as *pakɨma*. The vowels are different. So it matters that in the pairs in (1) – e.g. *pakma* vs *pakʊma* - [ʊ] is ***not*** epenthetic: this is sufficient to explain subject reactions whether or not the sequences in (1) are identified as well formed in English.

The long ISI used in the Korean experiments - 1500 ms – probably encouraged the subjects to encode words in memory in modified form relative to the stimulus perceived. It is known that longer ISI cause subjects to recast a stimulus in the more abstract terms of permissible phonological representations. Thus Werker and Logan (1985) find that different ISI categories relate to distinct modes of processing: acoustic, phonetic and phonemic. They show that short ISI (250ms) vs. medium (500ms) ISI produce a difference between acoustic vs. phonetic processing. Werker and Logan discuss ISI in the range of 1500ms and predict that listeners will perceive these entirely according to L1 phonological categories[6].

The 1500ms ISI used in Kabak and Idsardi's experiment have the property of making the results impossible to compare, for yet another reason, to those of Dupoux et al's Japanese experiments which were of just 500ms (reported for experiment 3 in Dupoux et al. 1999). Those results were striking because the ISI's were short enough to permit phonetic perception, under Werker and

[6] In the present case, we can surmise that bilinguals will be able to perceive stimuli separated by extra-long ISIs according to either L1 or L2 phonological categories.

Logan's classification. The results discussed here are less surprising: one might expect that long ISI result in a recoding of the stimuli in abstract terms, which filter out some phonetic details.

I have suggested here that the difference in discriminability between (1) and (2) stems from two factors, one of which would alone account for the results. First, Korean subjects could discriminate [gm]-[gʊm] but not [dʒm]-[dʒim] because the vowel is not epenthetic in the first pair, but is in the second. Second, clusters like [dʒm] required some form of repair, as they are ill-formed in both English (morpheme internally) and Korean (in all contexts): the status of clusters like [gm] is different, since they may be represented and stored as well formed English sequences[7]. Neither of these explanations invokes the difference that Kabak and Idsardi highlight between syllabic conditions (*strident-in-coda), which can cause perceptual epenthesis, and linear conditions (*stop-nasal) which can't. That distinction will have to be revised in any case, as it leaves unexplained the fact that Korean speakers adapt most non-native C_1C_2 clusters through epenthesis, even when the cluster violates a linear constraint: clusters like [dr], [tm] violate linear constraints but are resolved via epenthesis (*Fredrick* -> [p^{h}ured̠rik]; *litmus* -> [rit^{h}i̠mus̠]).

5. General issues

5.1. Syllable structure and segmental phonotactics

One general issue the authors raise is the relevance of syllables to the analysis of segmental phonotactics: the Korean experiment is presented as an argument in favor of syllable-based perception, along the lines of Mehler et al.'s (1981) earlier proposals. Kabak and Idsardi assume without argument that Korean [tʃ], [dʒ] are subject to a syllabic constraint which bans strident codas. In this section, I set aside the issues of interpretation raised earlier in order to address the question of choice between linear and syllabic constraints.

[7] The fact that English subjects did not repair the affricate-C sequences is clearly related to their greater experience with word-internal heteromorphemic sequences like the [dʒm] of *judgment*.

Earlier work (Chomsky and Halle 1968; Steriade 1999a,b, 2000; Blevins 2002) notes that syllable divisions are always predictable from a combination of segmental and morphemic information, so any syllabic constraint can be recast in non-syllabic, linear terms. The phonotactic conditions penalizing [gm] and [dʒm]-type clusters in Korean can both be stated linearly: a syllabic constraint like *Strident-in-coda corresponds, in Korean, to the linear statement "If C is strident then it's pre-vocalic", which exacts a penalty for every instance where the antecedent is true and the consequent is false. There is no formal defect in the latter statement and no loss of generality relative to the syllabic formulation, so no formal consideration tells us that the constraint targeting Korean stridents refers to onsets rather than to the pre-vocalic position. Likewise the constraint against Japanese "codas" (Ito 1987) is empirically equivalent to one requiring that consonantal place features occur pre-vocalically.

The same point can be made about the speech perception strategies Kabak and Idsardi discuss (p.32): "Hearing stridency can be said to make Korean listeners place the strident segment in onset position, which automatically evokes perceptual epenthesis". Why not "Hearing stridency can be said to make Korean listeners place the strident segment *in pre-vocalic* position, which automatically evokes perceptual epenthesis"?

While in Korean and Japanese the syllabic and linear constraints are descriptively equivalent, there are vast empirical evidences between them when we turn to the wider typology of consonantal phonotactics. Here the predictions of syllabic and linear or cue-based approaches diverge significantly (Blevins 2002, Côté 2001, Jun 2004, Kotchetov 1999, Steriade 1999a,b, 2000). To cite just one difference, in cases where the onset position does not happen to coincide with the position where a consonantal contrast is best cued – e.g. the case of the anteriority in apicals (Steriade 2000) – only the cue-based account provides a unified and predictive account of the phonotactic typology.

We have seen that there is no formal reason to think that the distribution of Korean and Japanese consonants is regulated by syllabic constraints, and that there are empirical grounds to expect the right constraints to be unrelated to syllables. Thus it's unclear that the results of Kabak and

Idsardi's study could shed light on the syllabic nature of perceptual illusions, even if their interpretation of the experimental results had not been open to question.

5.2.Underspecification in lexical access

The authors compared perception of [kC] to that of [gC] to determine whether voiced codas, which are impermissible in Korean, had a distinct ability to trigger perceptual epenthesis. Although the A' values of [gC]-[gVC] pairs are consistently lower, this does not appear to be a significant effect. The explanation proposed for this unexpected lack of perceptual epenthesis is that voicing is a redundant feature in Korean, and that voicing mismatches between stimulus and lexical entries are correspondingly ignored.

But the decision to call voicing a redundant feature in Korean remains unmotivated: the contrast between lax, tense and aspirated stops could be one between voiced, aspirated, and glottalized stops (with initial devoicing for lax stops). On this second interpretation voicing *is* distinctive, and the fact that Korean subjects did not perceive epenthetic vowels after the illegal voiced codas becomes problematic for the syllabic perception scenario advanced.

Even if voicing was unambiguously non-contrastive, we should note that the issue of contrast arises only internally to Kabak and Idsardi's interpretation of their data and vanishes when we adopt the scenario proposed in section 4: discrimination of, say, [gt]-[gʊt] is relatively robust not because [g] is just redundantly voiced, but because [ʊ] is not the epenthetic vowel of Korean. The results then do not speak to the role of contrast and underspecification in speech perception.

5.3. The P-Map

Kabak and Idsardi also discuss the P-map hypothesis (Steriade 2000, in press), a conjecture about the link between similarity rankings and the structure of the faithfulness component in an Optimality Theoretic grammar. The starting point for the P-Map is the observation that the choice of repair strategy for phonotactic violations is frequently predictable from the relative similarity of different input-output pairs. When the same phonotactic can be satisfied equally by two changes of the input, the typical choice is in favor of the less salient change. As a concrete

example, Korean speakers report that the pair [tm]-[t̬m] is intuitively more similar than [tm]-[mm]; the Korean rendition of English *litmus* is [lit^{h}t̬mus̬], not [limmus̬]. Both modifications -[tm]->[t̬m] and [tm]->[mm] succeed in satisfying all phonotactics violated by /tm/, but epenthesis preserves more of the input.

The precise content of the P-Map hypothesis is that a known similarity ranking – e.g. knowing that [tm]-[t̬m] is more similar than [tm]-[mm] - provides a default ranking of faithfulness constraints relative to each other, with the faithfulness constraints prohibiting greater input-output disparities ranked higher than those prohibiting lesser disparities. Then the primary respect in which relative similarity matters to phonology is in choosing between alternate modifications of the same string that would fix the same phonotactic violation, e.g. the choice between the /tm/->[t̬m] and the /tm/->[mm] repair of input /tm/. The P-map doesn't claim that perceptibility determines when a phonotactic will be satisfied or violated, or whether it will trigger any input modification at all: that would require a distinct hypothesis about the ranking of faithfulness constraints relative to competing phonotactics- not relative to competing faithfulness constraints – and it's not clear how knowledge of relative similarity could be linked to that idea. (A speculation to that effect is briefly discussed in Steriade 2000.)

It is perhaps this second hypothesis that Kabak and Idsardi have in mind when they discuss the relation between discriminability and input modification on p. 34 in connection with the P-Map. They discuss the claim, attributed to the P-Map, that an input and its phonotactically motivated output modification could not form a discriminable pair: such a claim is apparently contradicted by the Korean evidence, as /km/ can become [ŋm] in Korean and yet the [km]-[ŋm] pairs are discriminable by their Korean subjects.

In order to examine the P-Map's real predictions for this case, we need think about possible modifications of an input in relation to the constraint that triggers these modifications. So, regarding the change of Korean /km/ to [ŋm], suppose the constraint violated by /km/ which triggers this change is *$C_{[-son]}C_{[+son]}$. Now we need to think what options we have, in principle, to bring /km/ in compliance with this constraint. These options are limited to deletion of either C,

insertion of a vowel, and changes in the sonority of either C. Let's call these the set of relevant repairs. Each member of the set corresponds to a faithfulness constraint being violated[8]: varieties of MAX C correspond to different C-deletions, varieties of DEP V to different V-insertions, and Ident [±sonorant] to relevant feature changes. So the set of relevant repairs corresponds to a set of relevant faithfulness conditions. Now suppose that each member of the set of relevant repairs yields an input-output pair that exceeds some absolute discriminability threshold: so no matter how we fix /km/, the resulting input-output disparity is substantial, as smaller changes don't satisfy the phonotactic. In such a case, as noted earlier, the P-Map doesn't predict that /km/ will remain unmodified: to predict that, it would have to be a hypothesis about rankings between faithfulness and phonotactics. The P-Map does however predict one thing: if one member of the set of relevant repairs, R1, yields a more discriminable input-output pair than some other member, R2, then R1 will *not* be chosen[9]. This is because the P-Map uses relative similarity to project default faithfulness rankings, so the faithfulness constraint violated by the more salient modification, R1, will outrank the faitfulness constraint violated by R2. The actual prediction of the P-Map is then not about which phonotactics will be satisfied but about which repairs, from a larger set of relevant repairs, will *not* be resorted to.

A concrete prediction along these lines is this: if the pair /km/-[ŋm] is substantially less discriminable than /km/-[kɨm] then /km/ could not be modified as [kɨm]. Conversely, if /km/-[ŋm] is *more* discriminable than /km/-[kɨm], then /km/ could not be repaired as [ŋm].

These predictions are yet to be tested and the connection between specific discriminability differences and constraint rankings are yet to be explored.

[8] See McCarthy and Prince 1995 for an introduction to these constraints.
[9] Unless a conflicting phonotactic makes R2 inapplicable and leaves only R1 in the set of relevant repairs.

References

Blevins, Juliette. 2003. The Independent Nature of Phonotactic Constraints: An Alternative to Syllable-Based Approaches. In Caroline Féry and Ruben van de Vijver, Eds., The Syllable in Optimality Theory. Cambridge: Cambridge University Press, pp. 375-404.

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

Côté, Marie-Hélène. 2001. Consonant cluster phonotactics: a perceptual approach. PhD dissertation, MIT.

Dehaene-Lambertz, G., Dupoux, E., & Gout, A. 2000. Electrophysiological correlates of phonological processing: a cross-linguistic study. Journal of Cognitive Neuroscience, 12, 635-647.

Dupoux, E., K. Kakehi, Y. Hirose, 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.

Dupoux, E., Pallier, C., Kakehi, K., & Mehler, J. 2001. New evidence for prelexical phonological processing in word recognition. Language and Cognitive Processes, 5(16), 491-505

Dupoux, E. and S. Pepperkamp. 2002. The Phonetic Filter Hypothesis: How Phonology Impacts Speech Perception and Vice-Versa. presented at the Conference on Contrast in Phonology, University of Toronto.

Ito, Junko. 1988. *Syllable Theory in Prosodic Phonology*, Garland Publishing, New York.

Jun, Jongho. 2004. Place assimilation. In B. Hayes, R. Kirchner and D. Steriade (eds.) Phonetically Based Phonology. Cambridge University Press.

Jun, Sang-Beom. 2005. *Yeongeo Eumseonghak Kaelon.* (Introduction to English Phonetics). Seoul: Eulyu Mumhwasa.

Kabak, Baris and Idsardi, William (in press) Perceptual distortions in the adaptation of English consonant clusters: syllable structure or consonantal contact constraints. Language and Speech.

Kotchetov, Alexei. 1999. Constraints on the distribution of palatalized stops: Evidence for Licensing by Cue. In: Proceedings of North-Eastern Linguistic Society, Pius Tamanji, Masako Hirotani, and Nancy Hall (eds.). 167-181.

Mehler, J. J. Y. Dommergues, U. Frauenfelder, and J. Segui. 1981. The syllable's role in speech segmentation. Journal of Verbal Learning and Verbal Behavior, vol. 20, pp. 298-305.

Pepperkamp, S. and E. Dupoux. 2003. Reinterpreting loanword adaptations: the role of perception. Proceedings of the 15th ICPhS, Barcelona

Steriade, D. 1999a. Alternatives to the syllabic interpretation of consonantal phonotactics, in O.Fujimura B.Joseph and B.Palek (eds.) Proceedings of the 1998 Linguistics and Phonetics Conference, The Karolinum Press, 205-242

Steriade, D. 1999b. Phonetics in phonology: the case of laryngeal neutralization. In: Matthew Gordon (ed.), Papers in Phonology 3, UCLA Working Papers in Linguistics, Vol. 2, pp. 25-146.

Steriade, D. 2000. Directional asymmetries in place assimilation: a perceptual account. in E.Hume and K.Johnson (eds.) Perception in Phonology, Academic Press.

Steriade, D. (in press) The Phonology of Perceptibility Effects: the P-map and its consequences for constraint organization in Sharon Inkelas and Kristin Hanson (eds.) On the nature of the word, MIT Press.

Werker, J. F. and J. Logan. 1985. Cross-language evidence for three factors in speech perception. *Perception and Psychophysics*