

What we know about what we've never heard: evidence from production

Jason A. Shaw (Yale) & Shigeto Kawahara (Keio)

INTRODUCTION: Perceptual illusions have been used to argue for universal phonological preferences for sonority sequencing (Berent, Steriade, Lennertz, & Vaknin, 2007); however, this interpretation has also been challenged, as speech perception behavior is conditioned as well by interactions between language-specific phonetic experience and the demands of the particular experimental task (Daland, Oh, & Davidson, to appear; Davidson & Shaw, 2012; Peperkamp, 2007). This paper addresses the issue from a novel perspective in a way uniquely afforded by the pattern of high vowel devoicing in Japanese.

BACKGROUND: Japanese high vowels devoice between two voiceless consonants (e.g. [ɸusoku] ‘shortage’); recent articulatory research has shown that devoiced [u] is also variably deleted in this context (Shaw & Kawahara, 2018). Moreover, rates of deletion vary across items in ways that may have a phonological basis—deletion was more likely when it resulted in a sonority fall across syllables, e.g., /jutaisei/ → [j.tai.sei], than when it resulted in a sonority plateau, e.g., /ɸusoku/ → [ɸ.so.ku]. This pattern could be explained by scalar constraints regulating the sonority profile of syllable contact (see, e.g., Gouskova, 2004). However, since Japanese does not otherwise allow consonant clusters across syllable boundaries, there is no basis in the input to learn the relevant constraint weightings. Evidence from the processes of vowel deletion itself is not very informative for the learner, since the acoustic consequences of vowel deletion are masked by devoicing. Thus, any preference for falling sonority across syllables must emerge with little to no evidence for the relevant constraint weightings. To the extent that variation in vowel deletion is conditioned by phonological constraints on sonority sequences, this case may constitute stronger evidence for (unlearned) preferences than we have so far obtained from speech perception. However, as the relevant data from Shaw & Kawahara are rather thin—just two items, /jutaisei/ and /ɸusoku/, which besides syllable contact profile, differ on other dimensions, including the identity of C1—we conducted an experiment to more directly test whether vowel deletion in devoicing environments is conditioned by the resulting sonority profile across syllables.

EXPERIMENT: Articulatory trajectories of words containing /u/ in voiced and closely matched devoiced contexts were recorded using Electromagnetic Articulography (EMA). Seven native speakers of Tokyo Japanese produced 10-15 repetitions of each item in Table 1. The items were selected to include contexts for which vowel deletion would result in falling sonority profiles, fricative-stop (FS) sequences, and contexts for which vowel deletion would result in sonority plateaus, fricative-fricative (FF) sequences. Words were produced in a carrier phrase /okee ___ to itte/ ‘Okay, say ___’. Whether the vowel was present or deleted was determined through a Bayesian classification of tongue dorsum (TD) trajectories. The classifier was trained on voiced vowel TD trajectories (vowel present scenario) and a linear interpolation between flanking segments (vowel deletion scenario).

Table 1. Stimulus materials

Fricative Stop (FS)		Fricative Fricative (FF)	
voiced	devoiced	voiced	devoiced
ɸudou	ɸuton	ɸuzoku	ɸusoku
ɸudan	ɸutan	ɸuzai	ɸusai
ɸuda	ɸuta	ɸuzakeru	ɸusagaru

RESULTS: Figure 1 shows the posterior probability of vowel deletion by item. For all items, the distribution of posterior probabilities is distinctly bimodal—there are probability masses around 0 and 1 with not much in between, indicating that speakers are producing largely categorical variation. Tokens for which the tongue dorsum trajectory in devoiced vowels matches the vowel trajectory in voiced vowels have deletion probabilities at or near 0; tokens for which the tongue dorsum trajectory in devoiced vowels followed a roughly linear trajectory from the /e/ in the carrier phrase to the non-high vowel following /u/ (i.e., the underlined portion of /e#ɸusoku/) have a

What we know about what we've never heard: evidence from production

Jason A. Shaw (Yale) & Shigeto Kawahara (Keio)

probability at or near 1. Notably, there are few tokens with probability in the region of .5, which would indicate a vowel that was phonetically reduced. The categorical nature of the variation replicates the results of Shaw and Kawahara (2018). However, frequency of deletion is similar across items, including items resulting in FF sonority plateaus across syllables (red) and items resulting in FS sonority falls (grey). The effect of syllable contact (FF vs. FS) as assessed through the comparison of linear mixed effects models (with subject and item as random variables) was not significant. However, more interesting patterns emerged as we examine each speaker's behavior (Figure 2). Two speakers (S01, S04) produced nearly all tokens of devoiced vowels with a full vowel target (little to no deletion); two speakers (S02, S03) deleted vowels variably but equally across FF and FS items—for these speakers, the status of the resulting nuclei as [+continuant] may have been the determining factor for deletion probabilities; two speakers (S05, S06) produced the emergence of the unmarked pattern—they favored deletion in FS over FF; finally, one speaker (S07) showed the opposite pattern, favoring deletion in FF over FS.

DISCUSSION: Investigating vowel deletion at the level of articulation in devoiced vowels offers a unique opportunity to observe the effects of latent phonological knowledge. Our results replicated a key finding of Shaw & Kawahara—that high vowel deletion in Japanese is categorical (but variable) and that it gives rise to consonant clusters, and extends on it by showing that high vowels in similar phonological contexts behave similarly in terms of deletion probabilities, implying that deletion probabilities may be phonologically-conditioned rather than lexically-conditioned. When it comes to the specific form of the latent phonological knowledge, factors predicting the likelihood of vowel deletion varied across individuals. Syllable contact was a significant factor for just two of five speakers who exhibited variable deletion. Rather than seeing a consistent pattern emerge from a lack of evidence, as we would expect from innate (universal) preferences, we observed variation that included the pattern predicted by syllable contact constraints but also others.

Figure 1: deletion probability by item

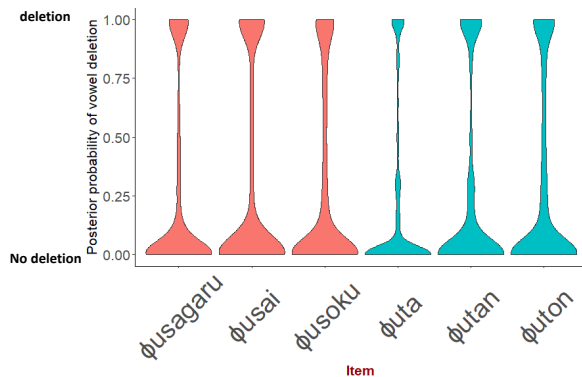
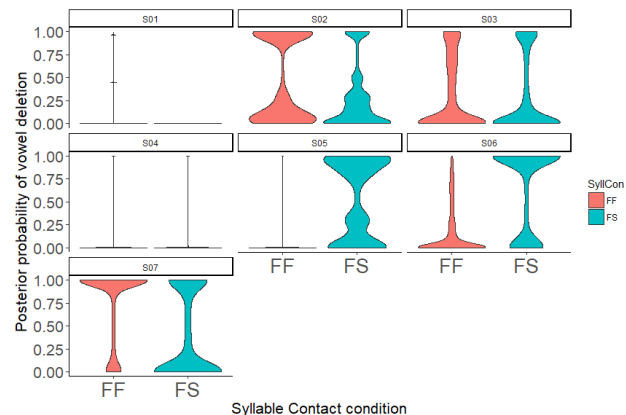


Figure 2: deletion probability by subject



References

- Berent, I., Steriade, D., Lennertz, T., & Vaknin, V. (2007). What we know about what we have never heard: Evidence from perceptual illusions. *Cognition*, 104, 591-630.
- Daland, R., Oh, M., & Davidson, L. (to appear). On the relation between online perception and loanword adaptation: Evidence from the perception of illicit word-medial clusters. *Natural Language & Linguistic Theory*.
- Davidson, L., & Shaw, J. A. (2012). Sources of illusion in consonant cluster perception. *Journal of Phonetics*, 40(2), 234-248.
- Gouskova, M. (2004). Relational hierarchies in Optimality Theory: the case of syllable contact. *Phonology*, 21(2), 201-250.
- Peperkamp, S. (2007). Do we have innate knowledge about phonological markedness? Comments on Berent, Steriade, Lennertz, and Vaknin. *Cognition*, 104, 631-637.
- Shaw, J. A., & Kawahara, S. (2018). The lingual articulation of devoiced /u/ in Tokyo Japanese. *Journal of Phonetics*, 66, 100-119. doi:<https://doi.org/10.1016/j.wocn.2017.09.007>