Cross-linguistic word-medial stop lenition: A Functional PCA approach

Seung Suk Lee [1,*,**], *Morgan Sonderegger* [1,*,**], *Meghan Clayards* [1,*,**]

[1] McGill University, Canada

seung.suk.lee@mcgill.ca, morgan.sonderegger@mcgill.ca, meghan.clayards@mcgill.ca

Abstract

Stop lenition between vowels is pervasive, yet fundamental questions remain: to what extent is lenition the same phenomenon across languages, and does it serve a common function? The *continuity lenition* hypothesis proposes one: weakened stops signal prosodic continuation and aid word segmentation. Previous work supports this idea but has examined few languages with differing methods, limiting cross-linguistic generalizability. We investigate word-medial continuity lenition in large speech corpora across nine languages, using Functional PCA of intensity contours to quantify the degree of lenition. We find a near-universal tendency for word-medial stops to be weaker, with voiced stops most consistent across languages. Exceptions reflect language-specific phonology or corpus artifacts, supporting continuity lenition as a cross-linguistic segmentation cue. Languages vary in whether weakening affects duration and/or intensity, suggesting listeners must learn cue patterns in their language.

Index Terms: corpus phonetics, lenition, functional PCA, crosslinguistic phonetics, stop consonants, prosody

1. Introduction

Stop lenition is an umbrella term for changes in the pronunciation of stop consonants where consonants are pronounced as 'weaker' [1]. Stop consonants are prototypically produced with an abrupt decrease in intensity, forming a 'dip' as the articulators create a complete closure that is sustained for a period. When stops are lenited (e.g., pronounced as a fricative or an approximant), the dip in the intensity contour is smaller and shallower due to an incomplete closure, and the stop closure is shorter [2]. While lenition is the most common type of variation in how stop consonants are pronounced [1, 3, 4, 5], fundamental questions about it remain—such as how stops vary across languages in the two common acoustic correlates of lenition.

Cross-linguistically, stops tend to be weakened between two vowels within a prosodic constituent (e.g., a word or a phrase) [1, 3, 4], and word-medial stops are shorter and more lenited [2, 6, 7, 8]. This weakening effect of word position is termed *continuity lenition* in [9], to convey the hypothesis that weakened stops signal the continuation of the ongoing prosodic constituent for listeners [2, 6, 7, 8, 10, 11].

Continuity lenition is closely related to *domain-initial strengthening*, another cross-linguistic tendency, for stronger and longer consonants at the beginnings of larger prosodic constituents [12, 13]. This typological tendancy for word-initial consonants to be longer than word-medial ones was recently confirmed in 43 out of 51 typologically diverse and understudied languages [14]. Both phenomena have been argued to serve as universal speech segmentation cues [15, 16], in line with perception experiments (e.g., [17, 18]) showing that listeners are sensitive to prosodically conditioned consonant variation.

While there seems to be a strong cross-linguistic tendency for stops to vary systematically by prosodic position, languages vary in the size of this prosodic boundary effect, its consistency across stop types (e.g., nasal vs. oral and voiceless vs. voiced), and in how consonantal strength is acoustically instantiated [8, 15, 19, 20, 21, 22, 23]. Previous studies have found that the acoustic duration of nasal stops is shorter at the beginning of a large prosodic constituent (Intonational Phrase) in American English [15, 21] and Seoul Korean [22, 23], while [8] found that Seoul Korean phrase-initial nasals show a more abrupt change in intensity contour than phrase-medial nasals. [21, 23] have argued these effects reflect nasals becoming more 'consonantal' domain-initially and weaker and more vowel-like domain-medially, consistent with the continuity lenition hypothesis.

Previous studies hint at a cross-linguistic universal of continuity lenition, with differences in its realization by language and by stop type. However, previous studies have examined only a handful of languages, typically with lab speech and differing operationalizations of lenition, making it difficult to assess how much different results reflect genuine language differences versus methodological differences. In line with other recent work in cross-linguistic corpus phonetics [5, 14, 24, 25], we address this using a unified methodology applied to large force-aligned corpora of read speech [26, 27] and spontaneous speech [28] across nine languages. To quantify the degree of stop lenition, we propose an approach using Functional Principal Component Analysis [29, 30] on intensity contours of intervocalic stops. We plan to expand to more languages in the future.

We investigate the following research questions: (1) Across languages, how consistent is the word-medial continuity lenition? (2) How consistent is it across stop types (voiceless/voiced oral stops, nasal stops)? We hypothesize that all word-medial stops are weaker than word-initial ones across all languages, but that languages and stop types may vary in *how* word-medial stops are weaker: some may be shorter but not more lenited in terms of the intensity transition (henceforth, just 'lenited'), while others may be more lenited but not shorter. Given previous findings, we expect languages to vary more for nasal [15, 21, 22, 23] or voiceless stops [20] than for voiced stops.

2. Methods

Our data come from three speech corpora. GlobalPhone [26] contains roughly 20 hours of read speech from 100 speakers per

[*] These authors contributed equally.
[**] indicates the corresponding author.

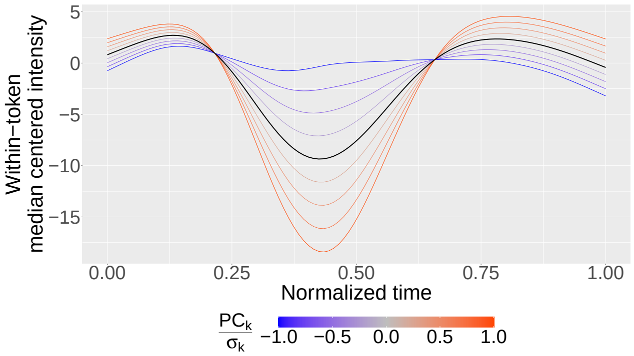

(a) *Average within-token median-centered intensity contour for /VCV/ (black line) ± 0.25 SD of PC1 (colored lines), across normalized time.*

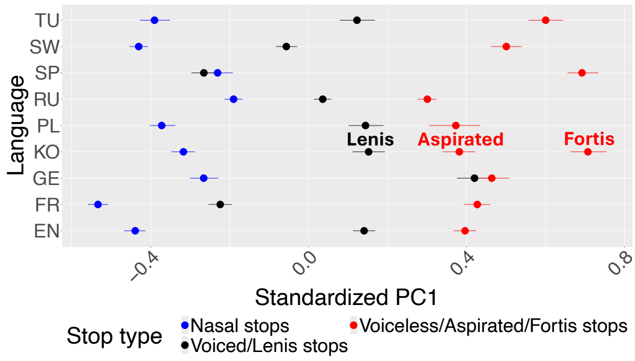

(b) *Model-estimated marginal median standardized PC1 for nasal, voiced, and voiceless stops across languages, from PC1 models. Dots and whiskers show median values and 95% highest posterior density intervals. Seoul Korean stop categories grouped as described in Sec. 2*

language; we used data from 7 languages (French, German, Polish, Russian, Spanish, Swedish, and Turkish). LibriSpeech [27] contains roughly 1000 hours of read speech from 800 American English speakers; we used a ∼20-hour subset (99 speakers). Both corpora were force-aligned with the Montreal Forced Aligner [31]. The Seoul Corpus [28] contains approximately 24 hours of spontaneous speech from 40 Seoul Korean speakers; the original transcription files were force-aligned as described in [28]. We used PolyglotDB [32] to extract acoustic measurements and query relevant tokens, finding 1.47m /V_1CV_2/ instances across 9 languages, where C was a nasal or oral stop varying in laryngeal categories (lenis/aspirated/fortis for Seoul Korean, voiced/voiceless for others) at a word-initial or word-medial syllable onset.

To operationalize the degree of lenition, we measured intensity at 10-ms intervals for each /V_1CV_2/ sequence. Previous studies have operationalized the degree of lenition using various aspects of the intensity contour, such as steepness of fall from the preceding vowel, magnitude of fall, and steepness of rise to the following vowel [2, 6, 7, 8, 11]. These differences complicate comparison of the size of the word boundary effect across languages. In addition, these measures assume the stop intensity forms a clear dip below neighboring vowels, which is not always the case for extremely lenited tokens; this has led some studies to discard up to 10% of tokens [2, 6, 8], potentially the most lenited tokens of interest. Our FPCA-based approach does not rely on finding such extrema and thus avoids this limitation.

Using FPCA, we identified the primary dimension of variation in intensity contours of intervocalic stops. Since our data included nasal, voiced, and voiceless intervocalic stops, which differ in intensity (nasal > voiced > voiceless), we hypothesized that the first principal component (PC) would capture relevant characteristics of the intensity contour: how large the intensity 'dip' was and how steep the transitions were. As a first step, we registered consonant boundaries as landmarks to time-normalize and smooth the intensity contours [30]. Intensity contours were median-centered within each token to parameterize intensity change relative to neighboring vowels. We performed FPCA on the intensity contours pooled across 9 languages, and found that the first PC captured both the size and slope of the intensity transition (as discussed further below). We quantified each stop token using its weight for this PC.

For statistical analyses, we fit two Bayesian mixed-effects linear regression models per language—one for the intensity PC (henceforth 'PC1 model') and one for the log-transformed phone duration (henceforth 'duration model')—using Stan/brms [33, 34]. We used flat priors for fixed effects and weakly informative priors for random effects (Student-*t* priors for variances, LKJ(1.5) priors for correlations). The posterior distribution of the model was sampled using Hamiltonian MCMC (4 chains, 4000 iterations/chain, 50% warm-up). Full model specifications and code are available at [35].

We used the models to test whether stop types (nasal, voiced, voiceless) varied systematically by word position (initial vs. medial) in intensity and duration. For Seoul Korean, stop type had four levels (nasal, lenis, aspirated, fortis). In presenting our results we group fortis/aspirated stops with voiceless and lenis with voiced stops of other languages (see [8, 36]).

Both models included major covariates of intensity and duration, including speech rate, place of articulation, neighboring vowel height, and syllable predictability [1, 2, 6, 7, 8, 10, 11]. Continuous predictors, as well as PC1 and log-transformed phone duration, were standardized within each language (centering and dividing by 2 SD), so that word-position effects are on the same scale, allowing direct comparison across languages. Models included three types of random effects: by-speaker, by-word indexed with syllable position (e.g., /p/ vs. /t/ in 'pata'), and by VCV sequence (e.g., /ata/ in 'pata'). We report models with random slopes for predictors of interest [37]: word-position and stop type.[1]

Word stress is known to condition stop realization ('stronger' articulation in stressed syllables) [38] and is confounded with word position in many languages (e.g. stress is predominantly word-initial in English [39] and German). To accurately estimate the word position effect we would ideally include stress as a covariate in the models. However, stress does not exist in some languages (e.g., Seoul Korean) or was not available in the pronunciation dictionaries in the corpora we used (GlobalPhone [26]). Instead, we conducted a supplementary analysis refitting our models for five languages for which partial stress information was available (English, German, Russian, Spanish, and Swedish[2]), to address the robustness of our answers to the research questions.

[1] We also fit models with maximal random-effect structure and found no qualitative differences.

[2] Swedish has lexical pitch accent rather than word stress. We use 'word stress' to refer to the lexically prominent syllable.

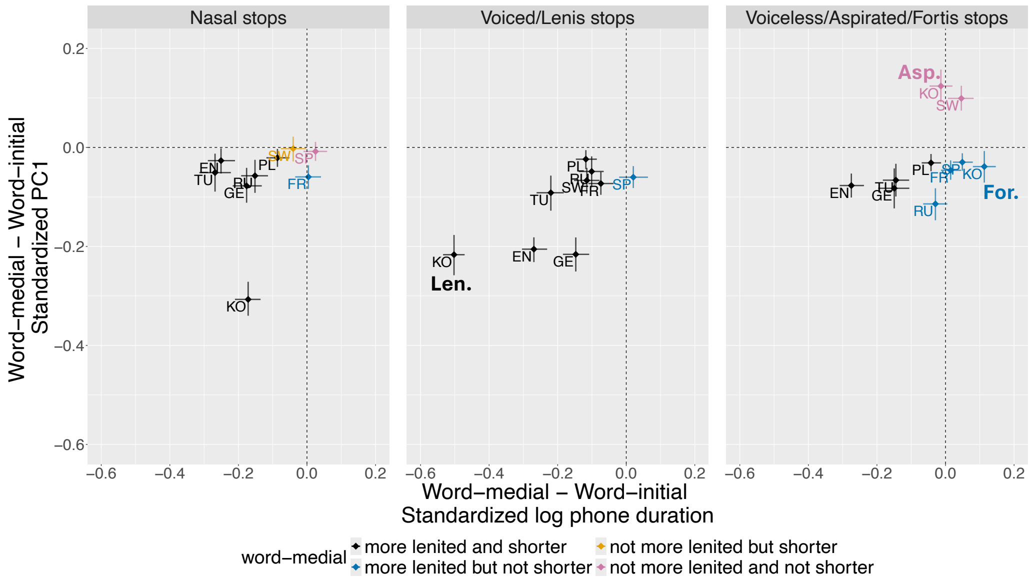

Figure 2: *Model-estimated differences in marginal median standardized log duration (x-axis) and standardized PC1 (y-axis) between word-medial and word-initial stops. Whiskers show 95% highest posterior density intervals. Color indicates whether word-medial stops are shorter (horizontal whiskers below 0) and/or more lenited (vertical whiskers below 0), or both.*

3. Results

Figure 1a presents the first PC, which accounts for 40.3% of the variation in the data. The higher the value of PC1, the more the intensity contour is shaped like the red curve relative to the blue curve. This PC captures both the size and steepness of the intensity transition. PC1 also distinguished stop types in each language, in the order expected under our interpretation—nasals < voiced stops < voiceless stops—as shown by the model predictions in Figure 1b (computed using the emmeans package [40]).

To address our main research question, we computed the difference between word-medial and word-initial positions for each stop type using the PC1 and duration models. The results are presented in Figure 2. We found that most cases supported our hypothesis: stops were 'weaker' word-medially in duration, PC1, or both. Almost all word-medial voiced (or Seoul Korean lenis) stops were both shorter and more lenited across all languages. In many cases, word-medial stops were more lenited but not shorter: French nasal stops, Spanish voiced stops, Russian, French, and Spanish voiceless stops, and Korean fortis stops. English and Swedish word-medial nasal stops were shorter but not more lenited. There were only three exceptions to our hypothesis: Spanish nasal, Korean aspirated, and Swedish voiceless stops. Spanish nasal stops showed no significant difference between the two word positions in either duration or PC1. Korean aspirated and Swedish voiceless were less lenited word-medially than word-initially. Korean fortis stops were also longer word-medially than word-initially, while Russian, French, and Spanish voiceless showed no significant difference between the two word positions in duration.

In our supplementary stress analysis, word stress was a significant predictor for both PC1 and duration across all five languages (English, German, Swedish, Russian, and Spanish): stops were shorter and more lenited as onsets of unstressed syllables, compared to stressed syllables (results available in [35]). Importantly, Figure 3 shows that including word stress in the models changed the sizes of word position effects but did not make the effects disappear, except for Swedish nasal stops. In this case, the word position effects on duration and PC1 were already small, and disappeared when word stress was included.

4. Discussion

We investigated the effect of word position on different stop types across nine languages. We first found that FPCA can be used to characterize the degree of 'dip' in intensity contours, distinguishing stop types in the expected order across languages. We then investigated how stops varied by word position, in each language. In most cases, word-medial stops were either shorter, more lenited, or both, compared to word-initial stops. Voiced stops were the most uniform across languages: consistently shorter and more lenited in 8 of 9 languages. Effects of word position on nasal stops varied more across languages, but no language showed a pattern counter to our hypothesis, i.e., word-medial lengthening or fortition. Effects of word position on voiceless stops also varied across languages, though word-medial stops were mostly more lenited.

One potential limitation of our results is that, despite using the same forced-alignment tool [31], there may be unrecognized cross-linguistic differences in alignment accuracy, particularly for highly lenited (or possibly elided) tokens.

Our results should also be interpreted with the caveat that 'word' is an approximation of prosodic structure: lexical word-medial stops are reliably prosodic word-medial, but word-initial stops may or may not be prosodic domain-initial. Some could be at the beginning of a larger prosodic constituent (strengthening our chance of finding a word-position effect), while others may be prosodic word-medial (weakening our chance), as in suffix-like orthographic words in Korean, which are always pronounced together with the preceding word [41]. This variance in what 'word-initial' means in our data introduces noise into our word-position results rather than systematically biasing them in one direction. Despite this, our results show a clear cross-linguistic tendency for word-medial stops to be weaker

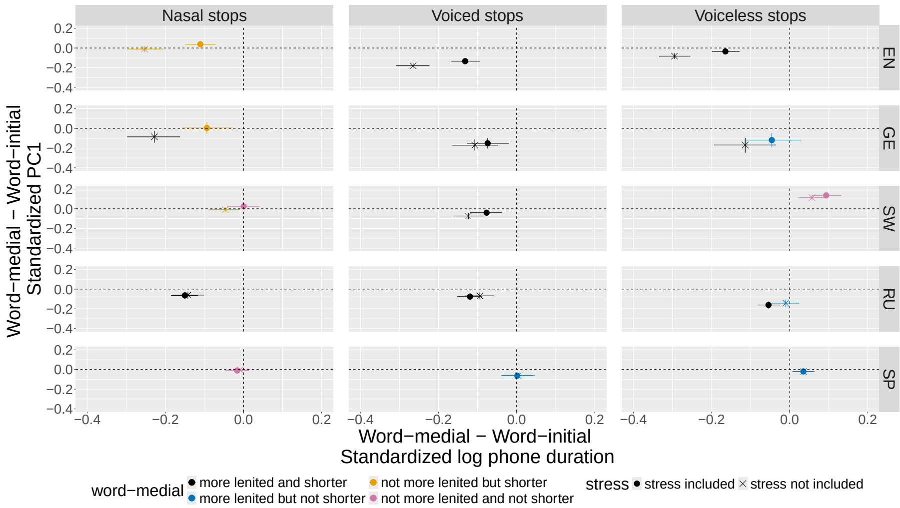

Figure 3: *Model-estimated differences in marginal median standardized log duration and standardized PC1 (word-medial minus word-initial). Crosses show estimates from models without word stress; circles show estimates from models including word stress. Whiskers show 95% highest posterior density intervals.*

than word-initial ones, suggesting the effect is robust.

The near-universality of word-medial weakening is further supported by the fact that the four exceptions to this pattern in Figure 2 each have explanations in language-specific phonology or aspects of the corpus used. The exception of Spanish voiced stops is consistent with [42], who finds that duration does not vary by prosodic position in Spanish and argues this is because intervocalic lenition is fully phonologized. The second exception is Swedish word-medial voiceless and nasal stops, which we believe is attributable to an artefact of GlobalPhone's pronunciation dictionary, which transcribes word-medial sonorants and voiceless stop geminates as singletons (consistent with the debated status of geminates in Swedish: [43]). Since geminates only occur word-medially in Swedish, this inflates the duration and intensity dip of word-medial stops relative to word-initial ones. Future work should correct for this issue.

Third, word-medial Korean fortis stops were longer than word-initial ones. This replicates [44] and extends their finding by showing that word-medial fortis stops are more lenited, despite being longer, than word-initial ones. Fourth, word-medial Korean aspirated stops were less lenited and not shorter than word-initial ones. This is surprising given previous findings of shorter VOT word-medially [16]. We speculate that this reflects vowel devoicing after word-initial aspirated stops in Seoul Korean [45]: devoiced vowels are lower in intensity, making the intensity transition from the stop to the following vowel more continuous, and thus lower in our PC1.

Our results extend previous work in two ways. First, while word-position effects on lenition, as measured by intensity dip, have been established for a handful of languages [2, 6, 7, 8, 11], we show these hold cross-linguistically, at least for voiced stops—supporting the idea that continuity lenition is a plausible universal cue to prosodic segmentation, consistent with [14]. This matters for listeners: acoustically 'strong' consonants aid listeners in detecting real [17] or learned artificial words [46], and in parsing prosodically disambiguated syntactic structures [18, 47, 48, 49]. Second, we show that languages vary in how duration and intensity are systematically affected by word position. This extends [14] by showing that duration alone is insufficient as a cue for segmentation, and suggesting that learning plays a role, as speakers need to learn how intensity and duration are used in their native language.

One example is Seoul Korean nasal stops, which stand out in the size of the word position effect on PC1, consistent with previous studies reporting an ongoing sound change of 'denasalization' of domain-initial nasal stops [18, 23, 50, 51]. However, Seoul Korean is the only spontaneous speech corpus in our sample; future work using a read speech corpus will be needed to determine whether its larger PC1 effect reflects genuine language-specific phonology or speech style differences.

Our supplementary analysis showed that word stress affects stop realization in the same way that the word boundary does, suggesting that the two effects are confounded. [38] has argued that word position effects for American English voiced stops are "mediated by stress". Including stress in our models affected the size of the word boundary effect, consistent with this view. However, we speculate that stress and word position may serve the same purpose as speech segmentation, to the extent that stressed syllables are word-initial. In English, German and Swedish, more than half of stressed syllables were word-initial while almost no unstressed syllables were, and these were also the languages showing the greatest shift in word position effect sizes when stress was included (shown in Figure 3). Listeners hearing a strongly pronounced stop in these languages thus receive a convergent cue to both word onset and stress.

5. Conclusion

We examined word-medial lenition across nine languages in large force-aligned corpora, quantifying lenition using an FPCA-based measure of intensity dip. We found a near-universal tendency for word-medial stops to be weaker than word-initial ones in duration and/or intensity, with voiced stops the most consistent across languages and exceptions largely attributable to language-specific phonology, supporting continuity lenition as a plausible cross-linguistic cue to prosodic segmentation. Future work will expand to more languages to test how universal these findings are.

6. Use of Generative AI

Generative AI was used to check grammar and compress prose to reach the page limit, without changing meaning.

7. Acknowledgments

This work was supported by Canada Research Chair #CRC-2023-00009 and NSERC grant #RGPIN-2023-04873 to Morgan Sonderegger, and SSHRC grant #435-2024-0996 to Meghan Clayards.

8. References

[1] R. M. Kirchner, "An Effort Based Approach to Consonant Lenition," Ph.D., University of California Los Angeles, 1998.

[2] J. Kingston, "Lenition," in *Selected Proceedings of the 3rd Conference on Laboratory Approaches to Spanish Phonology*, 2008.

[3] L. M. Lavoie, *Consonant Strength: Phonological Patterns and Phonetic Manifestations*, ser. Outstanding Dissertations in Linguistics. Florence: Taylor and Francis, 2001.

[4] N. Gurevich, "Lenition," in *Blackwell Companion in Phonology*, 2011, pp. 1–27.

[5] I. Vasilescu, Y. Wu, A. Jatteau, M. Adda-Decker, and L. Lamel, "Alternances de voisement et processus de lénition et de fortition: une étude automatisée de grands corpus en cinq langues romanes [voicing alternations in relation with lenition and fortition phenomena: an automated study of large corpora in five romance languages]," in *Traitement Automatique des Langues*, vol. 61, no. 1, 2020, pp. 11–36.

[6] J. Katz and G. Pitzanti, "The phonetics and phonology of lenition: A Campidanese Sardinian case study," *Laboratory Phonology: Journal of the Association for Laboratory Phonology*, vol. 10, no. 1, pp. 1–40, 2019.

[7] U. Cohen Priva and E. Gleason, "The causal structure of lenition: A case for the causal precedence of durational shortening," *Language*, vol. 96, no. 2, pp. 413–448, 2020.

[8] S. S. Lee, "Simultaneous encoding and decoding of laryngeal contrasts and prosodic structure in Seoul Korean," Doctoral Dissertation, University of Massachusetts Amherst, Amherst, Massachusetts, 2025.

[9] J. Katz, "Lenition, perception and neutralisation," *Phonology*, vol. 33, no. 1, pp. 43–85, 2016.

[10] J. Harris, "Grammar-internal and Grammar-external Assimilation," in *Proceedings of the 15th International Congress of Phonetic Sciences*. Barcelona, Spain: International Phonetic Association, 2003, pp. 1–4.

[11] T. Ennever, F. Meakins, and E. R. Round, "A replicable acoustic measure of lenition and the nature of variability in Gurindji stops," *Laboratory Phonology: Journal of the Association for Laboratory Phonology*, vol. 8, no. 1, pp. 1–32, 2017.

[12] T. Cho, "Prosodic Boundary Strengthening in the Phonetics–Prosody Interface," *Language and Linguistics Compass*, vol. 10, no. 3, pp. 120–141, 2016.

[13] E. Chodroff, "Phonetic Universals," *Annual Review of Linguistics*, vol. 11, no. 1, pp. 251–273, 2025.

[14] F. Blum, L. Paschen, R. Forkel, S. Fuchs, and F. Seifart, "Consonant lengthening marks the beginning of words across a diverse sample of languages," *Nature Human Behaviour*, vol. 8, no. 11, pp. 2127–2138, 2024.

[15] C. Fougeron and P. A. Keating, "Articulatory strengthening at edges of prosodic domains," *The Journal of the Acoustical Society of America*, vol. 101, no. 6, pp. 3728–3740, 1997.

[16] T. Cho and S.-A. Jun, "Domain-initial strengthening as enhancement of laryngeal features: Aerodynamic evidence from Korean," *UCLA Working Papers in Phonetics*, pp. 1–14, 2000.

[17] T. Cho, J. M. McQueen, and E. A. Cox, "Prosodically driven phonetic detail in speech processing: The case of domain-initial strengthening in English," *Journal of Phonetics*, vol. 35, no. 2, pp. 210–243, 2007.

[18] K. Yoo, "The production and perception of domain-initial strengthening in Seoul, Busan, and Ulsan Korean," Doctoral Dissertation, University of Cambridge, 2020.

[19] M. Hutin, Y. Wu, A. Jatteau, I. Vasilescu, L. Lamel, and M. Adda-Decker, "Synchronic fortition in five romance languages? a large corpus-based study of word-initial devoicing." in *Interspeech*, 2021, pp. 996–1000.

[20] P. Keating, T. Cho, C. Fougeron, and C.-S. Hsu, "Domain-initial articulatory strengthening in four languages," in *Laboratory Phonology VI: Phonetic interpretation*, J. Local, R. Ogden, and R. Temple, Eds. Cambridge: Cambridge University Press, 2004, pp. 143–161.

[21] T. Cho, D. Kim, and S. Kim, "Prosodically-conditioned fine-tuning of coarticulatory vowel nasalization in English," *Journal of Phonetics*, vol. 64, pp. 71–89, 2017.

[22] T. Cho and P. A. Keating, "Articulatory and acoustic studies on domain-initial strengthening in Korean," *Journal of Phonetics*, vol. 29, no. 2, pp. 155–190, 2001.

[23] J. Jang, S. Kim, and T. Cho, "Focus and boundary effects on coarticulatory vowel nasalization in Korean with implications for cross-linguistic similarities and differences," *The Journal of the Acoustical Society of America*, vol. 144, no. 1, pp. EL33–EL39, 2018.

[24] E. Chodroff, "Phonetic universals," *Annual Review of Linguistics*, vol. 11, no. 1, pp. 251–273, 2025.

[25] C. Ting, M. Clayards, M. Sonderegger, and M. McAuliffe, "The crosslinguistic distribution of vowel and consonant intrinsic f0 effects," *Language*, vol. 101, no. 1, pp. 1–36, 2025.

[26] T. Schultz, N. T. Vu, and T. Schlippe, "GlobalPhone: A multilingual text & speech database in 20 languages," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. Vancouver, BC, Canada: IEEE, 2013, pp. 8126–8130.

[27] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. South Brisbane, Queensland, Australia: IEEE, 2015, pp. 5206–5210.

[28] W. Yun, K. Yoon, S. Park, J. Lee, S. Cho, D. Kang, K. Byun, H. Hahn, and J. Kim, "The Korean Corpus of Spontaneous Speech," *Phonetics and Speech Sciences*, vol. 7, no. 2, pp. 103–109, 2015.

[29] J. O. Ramsay and B. W. Silverman, *Functional data analysis*, 2nd ed., ser. Springer series in statistics. New York, NY: Springer, 2005.

[30] M. Gubian, F. Torreira, and L. Boves, "Using Functional Data Analysis for investigating multidimensional dynamic phonetic contrasts," *Journal of Phonetics*, vol. 49, pp. 16–40, 2015.

[31] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal Forced Aligner: Trainable Text-Speech Alignment Using Kaldi," in *Proceedings of Interspeech 2017*, 2017, pp. 498–502.

[32] M. McAuliffe, E. Stengel-Eskin, M. Socolof, and M. Sonderegger, "Polyglot and Speech Corpus Tools: A System for Representing, Integrating, and Querying Speech Corpora," in *Interspeech 2017*. ISCA, 2017, pp. 3887–3891.

[33] B. Carpenter, A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell, "*Stan* : A Probabilistic Programming Language," *Journal of Statistical Software*, vol. 76, no. 1, 2017.

[34] P.-C. Bürkner, "**brms** : An *R* Package for Bayesian Multilevel Models Using *Stan*," *Journal of Statistical Software*, vol. 80, no. 1, 2017.

[35] S. S. Lee, M. Sonderegger, and M. Clayards, "Cross-linguistic word-medial stop lenition: A Functional PCA approach," 2026. [Online]. Available: https://doi.org/10.17605/OSF.IO/CE65H

[36] S.-A. Jun, "The Phonetics and Phonology of Korean Prosody," Doctoral dissertation, The Ohio State University, Columbus, Ohio, 1993.

[37] M. Sonderegger, *Regression modeling for linguistic data*. MIT Press, 2023.

[38] D. Bouavichith and L. Davidson, "Segmental and Prosodic Effects on Intervocalic Voiced Stop Reduction in Connected Speech," *Phonetica*, vol. 70, no. 3, pp. 182–206, 2013.

[39] A. Cutler and D. M. Carter, "The predominance of strong initial syllables in the english vocabulary," *Computer Speech & Language*, vol. 2, no. 3-4, pp. 133–142, 1987.

[40] R. V. Lenth, B. Bolker, P. Buerkner, I. Giné-Vázquez, M. Herve, M. Jung, J. Love, F. Miguez, H. Riebl, and H. Singmann, "emmeans: Estimated Marginal Means, aka Least-Squares Means," 2024. [Online]. Available: https://cran.r-project.org/web/packages/emmeans/index.html

[41] S.-A. Jun, "K-ToBI (Korean ToBI) Labelling Conventions (version 3.1)," 2000. [Online]. Available: https://linguistics.ucla.edu/people/jun/ktobi/k-tobi.html

[42] J. I. Hualde, M. Simonet, and M. Nadeu, "Consonant lenition and phonological recategorization," *Laboratory Phonology*, vol. 2, no. 2, 2011.

[43] Z. M. Hassan, "Gemination in Swedish & Arabic with a particular reference to the preceding vowel duration. An instrumental & comparative approach," in *Speech, Music and Hearing Quaterly Progress and Status Report (TMH-QPSR)*, vol. 44. Stockholm, Sweden: Department of Speech, Music and Hearing, KTH, 2002.

[44] Y. Yun, "Closure durations of Korean stops at three positions," *Phonetics and Speech Sciences*, vol. 14, no. 4, pp. 11–17, 2022.

[45] S.-A. Jun, M. E. Beckman, S. Niimi, and M. Tiede, "Electromyographic Evidence for a Gestural-Overlap Analysis of Vowel Devoicing in Korean," *UCLA Working Papers in Linguistics*, vol. 96, pp. 1–42, 1998.

[46] J. Katz and M. Fricke, "Auditory disruption improves word segmentation: A functional basis for lenition phenomena," *Glossa: a journal of general linguistics*, vol. 3, no. 1, 2018.

[47] D. R. Scott and A. Cutler, "Segmental phonology and the perception of syntactic structure," *Journal of Verbal Learning and Verbal Behavior*, vol. 23, no. 4, pp. 450–466, 1984.

[48] C. Kuzla, M. Ernestus, and H. Mitterer, "Compensation for assimilatory devoicing and prosodic structure in German fricative perception," in *Laboratory Phonology 10*, C. Fougeron, B. Kühnert, M. D'Imperio, and N. Vallée, Eds. DE GRUYTER MOUTON, 2010, pp. 731–758.

[49] H. Mitterer, S. Kim, and T. Cho, "The Role of Segmental Information in Syntactic Processing Through the Syntax–Prosody Interface," *Language and Speech*, vol. 64, no. 4, pp. 962–979, 2021.

[50] K. Yoshida, "Phonetic implementation of Korean denasalization and its variation related to prosody," *IULC Working Papers*, vol. 8, no. 1, 2008.

[51] J. Jang, J. Lee, J. Lee, S. Kim, and T. Cho, "Unveiling Denasalization as an Ongoing Sound Change: The Role of Prosody and Gender in Seoul Korean," *Language and Speech*, pp. 1–17, 2025.