Interspeech 2024
1-5 September 2024, Kos, Greece

Revisiting Pitch Jumps: F0 Ratio in Seoul Korean

Michaela Watkins[1], *Paul Boersma*[1], *Silke Hamann*[1]

[1]University of Amsterdam, The Netherlands

m.m.watkins@uva.nl, paul.boersma@uva.nl, silke.hamann@uva.nl

Abstract

Pitch tracking algorithms can show upward or downward jumps in F0 by one octave. These "octave jumps" are sometimes thought of as pitch-tracking "errors", in the sense that they constitute a "mistake" in the algorithm. Using Praat software, we discuss the point (which has been made before) that measured octave jumps often actually reflect genuine changes in periodicity and glottal-fold vibration. We illustrate this with the example of creaky voice in fortis stops in Seoul Korean. We argue (1) that when the goal is to capture periodicity or vocal-fold vibration, pitch-tracking algorithms capture F0 well, with pitch jumps possibly reflecting an important language-specific feature, and (2) that ignoring such jumps (due to assuming an error) could lead to misrepresentation of the properties of the language. To quantify these real F0 jumps, we introduce the notion of the "F0 ratio", which identifies potential F0 jumps and helps to chart the frequency of pitch jumps in a language.

Index Terms: Pitch tracking, creaky voice, Korean

1. Introduction

Creaky phonation is a type of voice quality in which (typically) a low and aperiodic or irregular fundamental frequency (F0) can be observed [1]. This type of phonation can be found across many different languages, and can constitute a contrast (cf. [2] for an overview). The pitch contour linked to creaky voice is susceptible to pitch jumps: the phenomena in which the F0 contour either increases or decreases abruptly by one octave (period doubling or period halving) [1].

Such octave jumps are sometimes considered pitch-tracking "errors", constituting a "mistake" in the pitch tracking algorithm. Where F0 is required for further investigation, this has resulted in strategies such as e.g. avoiding the entire section containing a jump, taking the mean overall F0 of the vowel, measuring the highest F0 observed, or calculating F0 at some point further into the vowel [3, 4, 5, 6, 7].

This paper argues that the above assumption is problematic. Talkin [8, p.500] already made the point that automated pitch-tracking algorithms do *not* necessarily make pitch tracking errors when detecting a jump, at least when the goal is to capture periodicity and glottal fold vibration. We support this view and provide an example, applying Praat [9] software to Standard Seoul Korean (SSK) fortis stops, to demonstrate how to detect and visualise F0 jumps that frequently occur at the beginning of the vowel following these stops. We also present a method of quickly assessing the number of pitch-jumping tokens by means of a simple "F0 ratio", and finally we discuss how there is no need to automatically ignore or avoid sections in the pitch curve that show a jump of any kind, while doing so could result in data misrepresentation.

Our aims for this paper are therefore the following:

- To discuss the notion that pitch jumps in pitch-tracking software should not automatically be considered errors, depending on what the goal of the measurement is.
- To present the example of Standard Seoul Korean, and demonstrate how ignoring tokens that exhibit pitch jumps would be detrimental to studying the language, given that such jumps occur in roughly 25% of fortis stops in SSK.
- To elaborate on ways of detecting and analysing such jumps, through both adjustment of the standard settings in Praat as well as through using the F0 ratio.

2. Pitch Tracking

2.1. What is pitch?

The term "pitch" can be defined narrowly to refer to a single auditory value for a perceptually matched height of a tone (or tone complex) [10]. In speech research one typically uses the term more loosely, but one can distinguish between at least three notions of pitch that correlate with the above narrow auditory definition [11]:

1) the frequency in which a wave shape repeats itself, i.e. periodicity or F0;
2) vocal-fold vibration frequency;
3) a more elusive and language-specific phonetic correlate of the intended linguistic intonation and/or tone contour.

Automated pitch-tracking algorithms in speech analysis software tend to be used for measuring any of these three. In the present paper we focus on the first two, as the paper concerns itself with vocal-fold vibration changes related to creaky voice (2), which influence the periodicity of the sound wave (1).

2.2. Pitch-tracking algorithms

As pitch-tracking algorithms almost invariably take as their input the speech waveform (rather than e.g. vocal-fold measurements), they define F0 as the "inverse of the smallest *true* period in the interval being analyzed" [8, p.497]. The present paper employs Praat [9] to measure this kind of periodicity, while interpreting sudden changes in measured F0 values as being caused by (slow or sudden) changes in vocal-fold vibration.

F0 estimation is challenging, because no speech signal is perfectly periodic: glottal periods are not always equally long and strong, and even if they are, the movements of the articulators change the shape of the waveform within consecutive periods. Additionally, F0 jumps of one octave up or down are not uncommon, and this is the focus of our investigation here.

10.21437/Interspeech.2024-1184

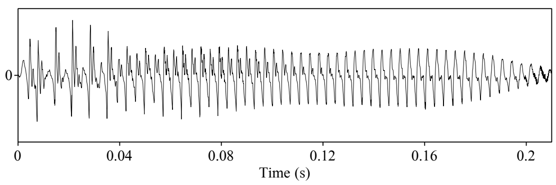

Figure 1: *A Korean recording of a vowel [a] (after a fortis stop)*

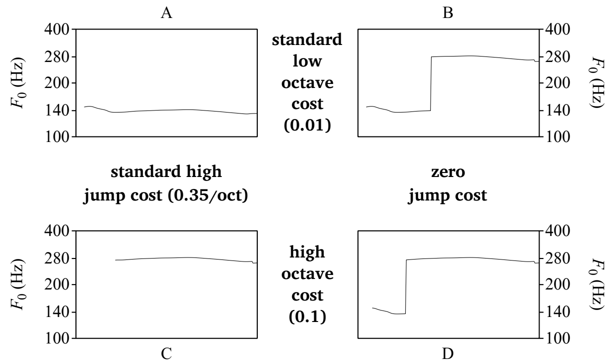

Figure 2: *Different options for settings to visualise the pitch jump of Figure 1 in Praat*

2.3. Ambiguous F0

Figure 1 shows a single speech recording, namely the vowel [a] (after a fortis stop that is not shown), produced by a female speaker of SSK in phrase-initial position. We can see that the pitch changes throughout the recording:

1) before 0.04 seconds, the F0 is approximately 140 Hz (because we see an unambiguous period of 7 ms);
2) after 0.10 seconds, the F0 is approximately 280 Hz (because we see an unambiguous period of 3.5 ms);
3) between 0.04 and 0.10 seconds the F0 is ambiguous: one can see peaks either every 3.5 ms or every 7 ms, depending on where one looks (the positive or negative peaks) and on how important one finds the (sometimes smaller, sometimes larger) differences between the heights of consecutive peaks.

As the pitch jump in Figure 1 reflects a genuine change in vocal-fold vibration (see §3), we want Praat to be able to detect the jump, i.e. to measure an F0 of 140 Hz in the beginning, one of 280 Hz at the end, with an upward jump somewhere in between. This is not entirely trivial, because the standard pitch settings in Praat do not accomplish this.

2.4. Pitch detection in Praat

The standard pitch settings in Praat have trouble detecting the pitch jump in Figure 1, for a good reason: these settings were optimized for finding the probably continuous linguistic intonation and tone contours that are intended by the speaker.

To explain what is going on, we describe how pitch detection works in Praat.

The pitch detection algorithm employed is called “raw autocorrelation” [12]. It works by establishing multiple F0 “candidates” for each local point in time, and then determining an optimal global F0 curve that runs through some of those candidates.

F0 candidates are determined by the time shifts over which the waveform is self-similar. In Figure 1, for instance, the F0 candidates are as follows:

1) Before 0.04 seconds, the waveform is self-similar over a time shift of 7 ms, so one F0 candidate is 140 Hz. However, the waveform is also similar over a time shift of 14 ms (i.e. two visibly determined periods), so another F0 candidate is 70 Hz. One could go on like this, with a time shift of 21 ms, giving an F0 candidate of 47 Hz, but such low candidates are ruled out by a setting called “pitch floor”, which is a frequency (e.g. 60 Hz) below which no F0 candidates are considered.
2) After 0.10 seconds, the waveform is self-similar over a time-shift of 3.5 ms (giving an F0 candidate of 280 Hz), a time-shift of 7 ms (giving an F0 candidate of 140 Hz), a time-shift of 10.5 ms (giving an F0 candidate of 93 Hz), and a time-shift of 14 ms (giving an F0 candidate of 70 Hz).
3) Between 0.04 and 0.10 seconds, the waveform is *somewhat* self-similar over time shifts of 3.5 ms and 10.5 ms, and *a bit more* self-similar over time shifts of 7 ms and 14 ms, giving strong F0 candidates of 70 and 140 Hz and somewhat weaker F0 candidates of 93 and 280 Hz.

Now, what is the best F0 candidate at each local point in time? Before 0.04 seconds, both 70 Hz and 140 Hz are perhaps equally strong candidates as far as self-similarity is concerned (technically speaking, both come with a self-similarity, i.e. relative autocorrelation peak, of almost 1), but it is clear that the visibly determined candidate of 140 Hz should win (cf. the word “smallest” in the quote in §2.2). This is established by the **octave cost** setting in Praat, which is standardly set at 0.01 and expresses a small preference for higher candidates (140 Hz is an octave higher than 70 Hz, so its strength is increased by 0.01 with respect to the 70-Hz candidate, in units of self-similarity). Likewise, the octave cost helps deciding that 280 Hz is the best local candidate for times after 0.10 seconds. Between 0.04 and 0.10 seconds, the octave cost’s preference for 280 Hz has to compete with the fact that the self-similarity is better for the 140-Hz candidate.

When determining the whole **global pitch curve**, an extra setting comes in, namely the **jump cost**, which is standardly set to 0.35/octave. This setting punishes the jump from 140 to 280 Hz that locally winning candidates would want to occur, and instead leads to a smooth curve hovering around 140 Hz (see Figure 2A). This path is chosen because 140 Hz is a candidate throughout the waveform (before 0.04 seconds, there is no 280-Hz candidate at all).

The path in Figure 2A may be appropriate for measuring intended (higher-level) intonation, because it starts early and is continuous, but it fails to detect the lower-level pitch jump that changes in the glottal-fold vibration have presumably caused.

2.5. Settings for detecting pitch jumps in Praat

Here we investigate what the optimal settings are for detecting the more local changes in periodicity that are caused by changes in vocal-fold vibration modes (as occur around creaky voice). The reader may think that simply setting the jump cost to zero might accomplish this, and indeed that will turn out to be the solution.

As we can see in **Figure 2**, if we leave the jump cost at its standard high setting, then the algorithm will detect either only the low F0 (140 Hz; Fig. 2A) or only the high F0 (280 Hz; Fig. 2C) depending on the octave cost setting. These two F0 curves have not heeded our wish to detect the jump.[1]

To detect the jump, we set the jump cost to zero, as in Figures 2B and 2D. The initial portion of the vowel shows a low overall F0 (140 Hz) before a large jump to a consistently stable higher F0 (280 Hz). The exact moment of the jump depends on the octave cost setting: with a high octave cost (0.1) it occurs around the beginning of the ambiguous part, whereas with the standard octave cost (0.01) the jump occurs around the middle of the ambiguous part. On the basis of such examples, we decide that good settings for detecting a pitch jump are a zero jump cost together with the standard octave cost of 0.01. This constitutes a minimal change from the standard settings in Praat: pitch jumps can be detected simply by setting the jump cost to zero and keeping all the other settings at their standard values.

3. Creaky Voice

An upward F0 jump, as we observed in our Korean example from Figure 1, is often associated with a transition from creaky voice to modal voice, while a downward F0 jump is often associated with the opposite transition. Why does a change from modal to creak and vice versa induce such jumps in the F0?

Creaky voice is typically defined as a type of voice quality in which the vocal folds vibrate with a low F0 and a long damping period [13]. Though there are different sub-types of creakiness, "prototypical" creaky voice shares a low F0 (low rate of vocal fold vibration), an irregular F0, and a constricted glottis (wherein vocal folds are close together) [1]. For listeners, a low as well as irregular F0 are sufficient for some kind of creaky voice to be perceived [14].

The irregularity or aperiodicity of the signal that is often linked to creak does not mean a random or unpredictable F0, but can cause problems for the pitch tracker to detect an F0 at all. This, however, is not the cause for a pitch jump. The jump is due to the genuine change in vocal fold vibration from a lower vibration rate in creaky voice to a higher one in modal voice, which is reflected in the acoustic signal as an ambiguous period (hence, an ambiguous F0), as seen in our example in Figure 1. One sub-type of creaky voice, the so-called "multiple-pulsing" [1, 14] or "period doubling" [15, 16], is characterised by exactly such a phase of ambiguous F0.

In the following, we briefly look at creaky voice in Mandarin Chinese, before we have a closer look at its manifestation in Standard Seoul Korean.

3.1. Mandarin Chinese

Creak in Mandarin is typically induced in low-F0 contexts and mostly seen on the low dipping tone (T3), but it also appears on other tones with low pitch [17, 18, 19], both in statements as well as in questions [20]. In Mandarin, creak is perceptually associated with low tone: extra low F0 improves identification accuracy of T3 compared to other tones, and irregular F0 hinders identification of T1 [19, 21]. With a raised F0, all tones, including T3, demonstrate less creak [18].

As for the *type* of creak, period doubling is most common across all four tones of Mandarin [15]. In perception, period-doubled creaky voice is interpreted as low and rough by Mandarin listeners [16], corresponding to the continuous low pitch in Fig. 2A of our Korean example.

We can conclude from these observations that creaky voice in Mandarin demonstrates that F0 and voice quality are intrinsically linked, and that an exclusion of pitch jumps from analysis would be problematic for this language.

[1]The first part of the pitch curve in Fig. 2C can be seen to have been analysed as voiceless, because there are no 280-Hz candidates there. The setting "voicing threshold", with a standard value of 0.45, locally militates against rejecting F0 candidates with a self-similarity over 0.45. This setting therefore counteracts the octave cost, in that it prefers the curve in Fig. 2A over the curve in Fig. 2C.

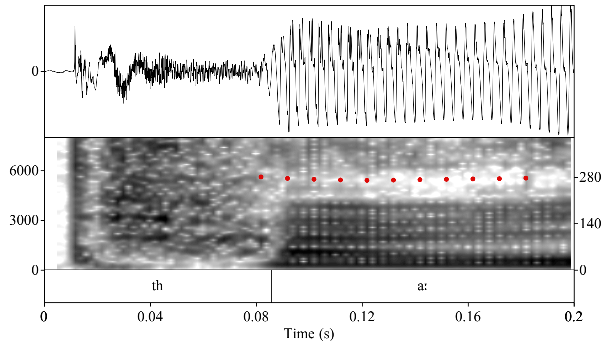

Figure 3: *Aspirated stop demonstrating no pitch jump*

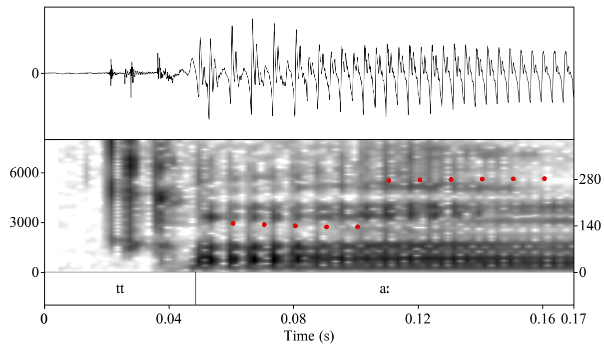

Figure 4: *Fortis stop demonstrating a pitch jump*

3.2. Standard Seoul Korean

Standard Seoul Korean (SSK) bears a typologically rare three-way laryngeal contrast in stops, as all three are voiceless in word-initial position. The discussion on how to distinguish the three is ongoing, and suggestions include (quasi-)tonogenesis [22, 23, 24, 25], as well as prosodic-dependent shifts in F0 and VOT [26, 27, 28, 29]. For this paper two of the three, namely fortis and aspirated, are selected to demonstrate the presence and absence of pitch-jumps. These two are related through similarities in "tensity" (i.e., a comparable high level of laryngeal activity [24, 30]), compared to the third stop type (lenis), which is sometimes described as "weak" [28, 29]. The two tense stops differ in terms of voice quality, with the fortis often being realized with creaky voice [4]. This creak can co-occur with a pitch jump, as we saw for the token in Figure 1.

In the discussion on Mandarin, we saw that low F0 cues creaky voice - this does not appear to be the case for SSK, as lenis stops consistently show lower F0 than fortis stops, without being creaky (e.g. [22, 24, 30]).

As can be seen in **Figure 3**, aspirated stops contain a long VOT with a relatively high F0 contour. Of interest here is the stability of this contour (visualised by red dots). For aspirated stops, a clear burst and no creak is heard (by the authors), with instead a long period of aspiration before vowel onset. **Figure 4** is a token of a fortis stop, the same that was used for Figure 1, exhibiting a pitch jump. For this token, creak is audible. (Both tokens were phrase-initial, and therefore did not undergo any co-articulatory influences.)

In contrast to Mandarin Chinese discussed above, creak and pitch jumps in SSK are less well studied, and different strategies have been observed when dealing with tokens of fortis stops that exhibit jumps in the following vowel: complete avoidance of the jump by considering the F0 only after the jump, usually someway into the vowel [3, 5, 4], by measuring the highest F0 [6], or by taking the mean F0 [7]. It is not yet known what perceptual impact creak has on Korean stops, nor whether the pitch jumps contribute to discriminating the three-way laryngeal contrast.

4. F0 Ratio

As a method of visualising how frequently tokens exhibit pitch jumps, we propose to use the "F0 ratio" with the settings mentioned previously in section 2: the jump cost set to 0 and the standard octave cost set to 0.01. To illustrate its use, we extracted all aspirated and fortis stops in word-initial position followed by [a] that were produced by 5 female speakers (between the ages of 16-18) in the *Korean Corpus of Spontaneous Speech* [31]. This resulted in a total of 270 tokens. The vowel following each stop was annotated in Praat [9], then median F0 values for the first and last third of the vowel were determined, and the F0 ratio was calculated using the formula in (1).

$$\text{F0 ratio} = \frac{\text{median F0 in last 1/3 of vowel}}{\text{median F0 in first 1/3 of vowel}} \quad (1)$$

The F0 ratios for both stop types in SKK are plotted in **Figure 5**. For the aspirated stops in the left half, the majority of tokens has a ratio of around 1, indicating that no pitch jump is present. Three stops have a higher ratio, around 2, indicating pitch doubling, and three a lower one, around 0.5, indicating pitch halving. We propose that F0 ratio values between 0.8 and 1.25 are suggestive of no pitch jumps present, while values roughly above 1.7 indicate pitch doubling, and values lying below 0.6 indicate pitch halving, as summarized in **Table 1**. These values may sound arbitrary, but the whole of Figure 5 indicates that F0 ratios in between the three ranges (i.e. between 0.6 and 0.8, or between 1.25 and 1.7) may be quite rare, so that our three-way binning of the continuous F0 ratios (halving, no jump, and doubling/tripling) seems legitimate.

Table 1: *How to bin the F0 ratios*

Ratio	Jump Type
$0.8-1.25$	No Jump
> 1.7	Doubling/Tripling
< 0.6	Halving

For the fortis stops in the right half of Figure 5, the majority of tokens cluster again around a ratio of 1, but now roughly 25% either show pitch doubling or halving (as defined in Table 1), which is an amount that does not seem to be attributable to other phonetic effects than creak. This supports again the notion that such pitch jumps are not measurement errors, but rather possibly an inherent phonetic correlate of a linguistic category or feature of the language in question.

The F0 ratio allows us to visualise tokens quickly and observe patterns in the data: if tokens form one cluster with few outliers, such as the aspirated stops here, this may indicate the absence of consistent pitch jumps. Individual stops may then be manually checked for phonetic effects such as e.g. coarticulation of preceding vowels in continuous speech or interference such as background noise. For the fortis stops, roughly 25% of tokens showed some kind of jump - this is enough to suggest that they deviate systematically, warranting further investigation. We conclude that the F0 ratio method allows us to select and assess tokens with a jump individually, rather than removing them from further analysis.

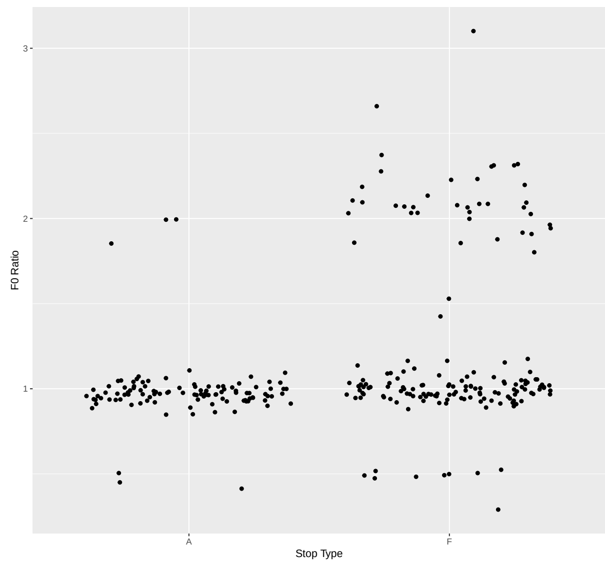

Figure 5: *Distribution of F0 Ratio in Seoul Korean vowels following tense stops. 'A' = Aspirated (left), 'F' = Fortis (right)*

5. Discussion

The present paper emphasizes (with [8]) that pitch jumps detected by pitch-tracking algorithms are often not measurement errors but instead reflect genuine sudden changes in F0 caused by (slow or sudden) changes in the mode of vocal-fold vibration, as often happens around creaky voice. To assist future research addressing pitch jumps, we suggest here using the F0 ratio as an initial step to provide an overview of the data in question – if no jumps are observed, then the researcher can proceed with further F0 analyses. Should jumps be observed via the F0 ratio, it is best to then manually assess the tokens with jumps before continuing with further F0-based calculations. Reliable F0 measurements are needed, for instance, for reliably computing measures such as H1−H2, which is a common method of evaluating voice quality in Korean [5, 7, 23, 32]. It should be noted that this study uses a single language, therefore future studies should incorporate more variety of languages (including non-tonal ones) to validate the robustness of this measure. Finally, the relevance of our findings for the *perception* of Standard Seoul Korean remains to be investigated. Creak is applicable to only one of the three stop types in SSK. It is not yet known whether creaky voice in SSK assists with the discrimination of the laryngeal contrast, and whether the pitch jumps that we observed are at all perceptible to the listener.

6. References

[1] P. Keating, M. Garellek, and J. Kreiman, "Acoustic properties of different kinds of creaky voice," in *Proceedings of the 18th International Congress of Phonetic Sciences*, 2015, pp. 0821–1.

[2] M. Gordon and P. Ladefoged, "Phonation types: a cross-linguistic overview," *Journal of Phonetics*, pp. 383–406, 2001.

[3] J. J. Holliday and E. J. Kong, "Dialectal variation in the acoustic correlates of Korean stops," in *Proceedings of the 17th International Congress of Phonetic Sciences*, 2011, pp. 878–881.

[4] Y. Kang, J. Schertz, and S. Han, "The phonology and phonetics of Korean stop laryngeal contrasts," in *The Cambridge handbook of Korean linguistics*. Cambridge University Press, 2022, pp. 215–247.

[5] Y. Kang, "Voice onset time merger and development of tonal contrast in Seoul Korean stops: A corpus study," *Journal of Phonetics*, vol. 45, pp. 76–90, 2014. [Online]. Available: http://dx.doi.org/10.1016/j.wocn.2014.03.005

[6] K.-H. Lee and M.-S. Jung, "Acoustic characteristic and perceptual cues for Korean stops," *Speech Sciences (음성과학)*, vol. 7, pp. 139–155, 2000.

[7] M. Oh and H. Yang, "The production of stops by Seoul and Yanbian Korean speakers," *Journal of the Korean Society of Speech Sciences*, vol. 5, pp. 185–193, 2013.

[8] D. Talkin, "A robust algorithm for pitch tracking (RAPT)," in *Speech Coding and Synthesis*, W. B. Kleijn and K. Paliwal, Eds. Elsevier Science, 1995, pp. 495–518.

[9] P. Boersma and D. Weenink, "Praat: Doing phonetics by computer," 2024, [computer program]. Version 6.4.06, retrieved 7 March 2024 from http://www.praat.org.

[10] B. C. Moore, B. R. Glasberg, and R. W. Peters, "Relative dominance of individual partials in determining the pitch of complex tones," *The Journal of the Acoustical Society of America*, vol. 77, pp. 1853–1860, 1985.

[11] P. Boersma, "Acoustic analysis," in *Research methods in linguistics*, R. Podesva and D. Sharma, Eds. Cambridge University Press, 2013, pp. 375–396.

[12] ——, "Accurate short-term analysis of the fundamental frequency and the harmonics-to-noise ratio of a sampled sound," in *Report of the Institute of Phonetic Sciences Amsterdam*, vol. 17, 1993, p. 97–110.

[13] L. Davidson, "The versatility of creaky phonation: Segmental, prosodic, and sociolinguistic uses in the world's languages," *Wiley Interdisciplinary Reviews: Cognitive Science*, vol. 12, p. e1547, 2021.

[14] P. A. Keating, M. Garellek, J. Kreiman, and Y. Chai, "Acoustic properties of subtypes of creaky voice," *The Journal of the Acoustical Society of America*, vol. 153, p. 297, 2023.

[15] Y. Huang, "Articulatory properties of period-doubled voice in Mandarin," in *Proceedings of the International Conference on Speech Prosody*. International Speech Communication Association, 2022, pp. 545–549.

[16] ——, "Perception and imitation of period doubling," in *Proceedings of the International Congress of Phonetic Sciences*, 2023, pp. 327–331.

[17] R. Cao, R. Wayland, and E. Kaan, "The role of creaky voice in Mandarin tone 2 and tone 3 perception," in *Proceedings of the 13th Annual Conference of the International Speech Communication Association*, vol. 1, 2012, pp. 426–429.

[18] J. Kuang, "Covariation between voice quality and pitch: Revisiting the case of Mandarin creaky voice," *The Journal of the Acoustical Society of America*, vol. 142, pp. 1693–1706, 2017.

[19] Y. Huang, "The role of creaky voice attributes in Mandarin tonal perception," in *Proceedings of the International Congress of Phonetic Sciences*, 2019, pp. 1465–1469.

[20] Y. Chai, "The source of creak in Mandarin utterances," in *Proceedings of the 19th International Congress of Phonetic Sciences*, 2019, pp. 1858–1862.

[21] Y. Huang, "Different attributes of creaky voice distinctly affect Mandarin tonal perception," *The Journal of the Acoustical Society of America*, vol. 147, pp. 1441–1458, 2020.

[22] H. Y. Bang, M. Sonderegger, Y. Kang, M. Clayards, and T. J. Yoon, "The emergence, progress, and impact of sound change in progress in Seoul Korean: Implications for mechanisms of tonogenesis," *Journal of Phonetics*, vol. 66, pp. 120–144, 2017. [Online]. Available: https://doi.org/10.1016/j.wocn.2017.09.005

[23] K.-H. Kang and S. G. Guion, "Clear speech production of Korean stops: Changing phonetic targets and enhancement strategies," *The Journal of the Acoustical Society of America*, vol. 124, pp. 3909–3917, 2008.

[24] D. Silva, "Acoustic evidence for the emergence of tonal contrast in contemporary Korean," *Phonology*, vol. 23, pp. 287–308, 2006.

[25] ——, "Variation in voice onset time for Korean stops," *Korean Linguistics*, vol. 13, pp. 1–16, 2006.

[26] T. Cho, "The phonetics-prosody interface and prosodic strengthening in Korean," in *The Cambridge handbook of Korean linguistics*. Cambridge University Press, 2022, pp. 248–293.

[27] J. Choi, S. Kim, and T. Cho, "An apparent-time study of an ongoing sound change in Seoul Korean: A prosodic account," *PLoS ONE*, vol. 15, pp. 1–29, 2020. [Online]. Available: http://dx.doi.org/10.1371/journal.pone.0240682

[28] S. A. Jun, "Influence of microprosody on macroprosody: a case of phrase initial strengthening," *UCLA Working Papers in Phonetics*, vol. 92, pp. 97–116, 1996.

[29] ——, "No tonogenesis in Korean, yet," in *Proceedings of the World Congress of Korean Language*, 2020, p. 10. [Online]. Available: https://linguistics.ucla.edu/people/jun/papers%20in%20pdf/49-Jun_2020%20World%20Congress%20of%20Korean%20Language.pdf

[30] C.-W. Kim, "On the autonomy of the tensity feature in stop classification (with special reference to Korean stops)," *WORD*, vol. 21, pp. 339–359, 1965.

[31] 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," *Journal of the Korean Society of Speech Sciences*, vol. 7, no. 2, p. 103–109, 2015.

[32] Y. Chai and M. Garellek, "On H1–H2 as an acoustic measure of linguistic phonation type," *The Journal of the Acoustical Society of America*, vol. 152, pp. 1856–1870, 9 2022.