

# The effect of different information sources on prosodic boundary perception

Bogdan Ludusan,<sup>1,a)</sup> Masahiro Morii,<sup>2</sup> Yasuyo Minagawa,<sup>3</sup> and Emmanuel Dupoux<sup>1</sup>

<sup>1</sup>Laboratoire de Sciences Cognitives et Psycholinguistique, EHESS/ENS-PSL/CNRS/INRIA, 24 rue Lhomond, Paris, 75005, France

<sup>2</sup>Global Research Center for Logic and Sensibility, Keio University, 2-15-45 Mita, Minato-ku, Tokyo, 108-8345, Japan

<sup>3</sup>Department of Psychology, Keio University, 4-1-1 Hiyoshi, Kohoku-ku, Yokohama, 223-8521, Japan

bogdan.ludusan@uni-bielefeld.de, masa.morii@gmail.com, myasuyo@bea.hi-ho.ne.jp, emmanuel.dupoux@gmail.com

**Abstract:** This study aims to quantify the effect of several information sources: acoustic, higher-level linguistic, and knowledge of the prosodic system of the language, on the perception of prosodic boundaries. An experiment with native and non-native participants investigating the identification of prosodic boundaries in Japanese was conducted. It revealed that non-native speakers as well as native speakers with access only to acoustic information can recognize boundaries better than chance level. However, knowledge of both the prosodic system and of higher-level information are required for a good boundary identification, each one having similar or higher importance than that of acoustic information. © 2021 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).

[Editor: Martin Cooke]

<https://doi.org/10.1121/10.0007150>

Received: 29 June 2021 Accepted: 22 October 2021 Published Online: 17 November 2021

## 1. Introduction

Prosody is an integral component of spoken language and the prosodic structure, by means of the edges of its different constituents—*prosodic boundaries*, has been shown to be used in various processes, from language acquisition in infants [e.g., word segmentation (Johnson *et al.*, 2014)] to language processing in adults [e.g., sentence parsing, see Cutler *et al.* (1997) for a review].

Listeners have at their disposal various knowledge sources to postulate an upcoming prosodic boundary, including higher-level linguistic information (language-specific, in the form of the lexical identity of the words, their meaning, or the word order used in the language as well as acoustic information, either language-independent (the actual values of the acoustic cues, as present in the speech signal) or language-specific (knowledge of the prosodic system of the language, e.g., the importance/weight of each cue in the marking of prosodic boundaries). Despite considerable interest in the scientific community, no previous studies have examined the role of all three information sources, nor did they attempt to quantify their effects.

Several works have looked at the role played by the language-independent acoustic component in boundary recognition by presenting stimuli to native and non-native participants [e.g., Carlson *et al.* (2005) and Kuo (2014)]. As both groups were able to predict the boundary strength, these studies concluded that mainly acoustic information is used by listeners to posit prosodic boundaries. Other studies have shown, instead, that also higher-level information plays an important role in the perception of prosodic boundaries. For instance, the results of Simon and Christodoulides (2016) in which stimuli were delocalized by means of re-synthesis, revealed that not only acoustic, but also syntactic cues are used for determining the existence of a prosodic boundary in running speech. Moreover, Cole *et al.* (2010) holds that syntactic information has the most important role in the perception of boundaries, followed by acoustic information. Further evidence that higher-level linguistic knowledge plays a role in the recognition of prosodic boundaries can also be seen from speech technology. Automatic systems for prosodic boundary detection reach a high performance when lexical, semantic, or syntactic knowledge is used for classification (Sloan *et al.*, 2019) or when these are added to an acoustic-only model (Ananthakrishnan and Narayanan, 2008).

Besides language-independent acoustic and higher-level linguistic information, also other knowledge sources may play a role in the perception of prosodic boundaries. For example, although boundaries tend to be marked by the same acoustic cues across languages, the weight given to each cue varies with the language (Vaissière, 1983). Thus, a person familiar with a language would use this (language-learned) acoustic information differently from a person who is not familiar with that language. Confirmation of this familiarity effect was found by Wakefield *et al.* (1974), in which non-

<sup>a)</sup> Author to whom correspondence should be addressed. Current address: Faculty of Linguistics and Literary Studies, Bielefeld University, Universitätsstraße 25, 33615 Bielefeld, Germany, ORCID: 0000-0002-2701-6569.

native participants previously exposed to the test language, had a higher recognition rate of prosodic boundaries, compared to a group with no exposure to the test language. Further evidence towards the usefulness of language-learned acoustic information can be seen in the results of Kuo (2014), with English-native listeners being able to discriminate the prosodic boundary strength of one-word Swedish stimuli, while Taiwanese listeners could not (English exhibiting more similarities to Swedish, than Mandarin).

The goal of this study is to better understand how these various information sources are used in the perception of prosodic boundaries and to quantify their effect. For this, we test both native and non-native speakers using Japanese materials. The Japanese prosodic organization features two levels of phrasing below that of the utterance, called intonation and accentual phrases, defined based on intonation and on the degree of perceived disjuncture (Venditti, 2005). The intonation phrase is the higher of the two levels and it marks a stronger degree of disjuncture than the accentual phrase. Tonally, accentual phrases contain at most one lexical pitch accent and exhibit a pitch rise on their first or second mora, followed by falling pitch towards the end of the phrase. A process of “downstep” occurs at this level, by which the pitch height of each phrase decreases if it was preceded by a lexically accented phrase. By contrast, at the intonation phrase level, the speaker sets a new pitch range at its beginning, which is independent from that of the previous intonation phrase. In terms of the acoustic cues marking prosodic boundaries, it has been shown that cues such as pause, nucleus duration or pitch reset are employed in Japanese (Vaissière, 1983), although they may be used differently across the two levels of prosodic phrasing [e.g., Ludusan *et al.* (2016)].

Here, we use a similar experimental paradigm to the one employed by Wakefield *et al.* (1974), by splicing pauses at prosodic boundary positions vs phrase-internally (as given by the annotation supplied with the corpus employed to extract the stimuli) and asking our subjects to judge which one of the two versions sounds more natural. Early psycho-acoustic studies [e.g., Wakefield *et al.* (1974)] have shown that listeners tend to judge stimuli having pauses inserted at prosodic boundary location as being more natural. Our approach is also similar to the one used by Simon and Christodoulides (2016), as it determines the existence/absence of a prosodic boundary, thus allowing us to compute an accuracy of the process. The performance of the recognition process would give us insights into the usefulness of prosodic boundary information for various linguistic processes (e.g., word segmentation in early language acquisition). Moreover, we test more conditions than previously employed in studies on prosodic boundary perception, giving us a more complete view of the process and also permitting us to perform methodological comparisons (e.g., of different de-lexicalization methods). We included both native [as in Carlson *et al.* (2005), Cole *et al.* (2010), and Simon and Christodoulides (2016)] and non-native [e.g., Carlson *et al.* (2005), Kuo (2014), and Wakefield *et al.* (1974)] participants listening to non-manipulated speech (containing all information sources), as well as to manipulated speech [either re-synthesized, similar to Simon and Christodoulides (2016) or filtered, such as Kuo (2014)].

## 2. Methods

An overview of the employed stimuli and the experimental procedure is given in this section, with a more detailed description of the methods, as well as the list of stimuli used in this study and an acoustic analysis of the latter being included as supplementary material.<sup>1</sup>

### 2.1 Stimuli

The Corpus of Spontaneous Japanese (Maekawa, 2003) was used for creating the experiment stimuli, as its core part has been annotated at the segmental and prosodic levels, including annotations for prosodic boundaries. We considered the academic sub-part of the corpus, and determined all inter-pausal units which ended with an intonation phrase (IP) boundary [equivalent to a level 3 break, according to the J-ToBI model (Venditti, 2005)], were preceded by an inter-pausal unit also ending with an IP boundary, had no internal pauses and had only one internal prosodic boundary, which was also an IP boundary. We then limited our search to inter-pausal units having a minimum of four syllables and a maximum duration of 2 s and the resulting items were vetted by two native speakers to contain a full meaning. The 80 items that passed this check were employed as stimuli in our experiment.

The first condition, *nonManip*, used the original, non-manipulated items. Each item had two versions created: one in which a 300 ms pause was spliced at the position on the internal IP boundary (further called *boundary* position) and another one in which an identical pause was added somewhere else in the item (called *non-boundary* position). We take as example one of the considered items, まず背景ですが: “First of all, about (the study) background,” which has an internal IP boundary between the words まず and 背景. The boundary position stimulus derived from it had a pause inserted between the words まず and 背景, while the non-boundary stimulus had an identical pause inserted between the words 背景 and です. The non-boundary position was balanced such that half of the stimuli had the pause introduced at a word boundary and half word-internal (syllable boundary).

In the second condition, *synthesized*, we removed the higher-level linguistic information by means of re-synthesis. The original stimuli were re-synthesized by replacing all vowels by the vowel /a/, all nasals by /n/, all stops by /t/, all fricatives by /s/ and all glides by /j/, employing appropriate voice models to the gender of the person that uttered

the original stimuli. The same procedure was then followed to introduce pauses at boundary and non-boundary positions, respectively.

The third condition, *filtered*, saw the original stimuli low-pass filtered with a cut-off frequency of 300 Hz and a smoothing of 100 Hz, thus effectively retaining only the tonal information. After filtering the signal, pauses were spliced in order to create the two versions of the stimuli, as performed for the previous two conditions.

### 2.2 Procedure

25 Japanese native speakers, having grown up in a monolingual family in the Kanto area and not having lived abroad for longer than three months, participated in each of the three conditions. They were tested in a sound-proof booth, using a computer running the PSYCHOPY software (Peirce, 2007). They were told they would be hearing sentences that have been artificially modified in two different ways and that they are supposed to judge which of the two versions sounds more natural. Based on previous research (Wakefield et al., 1974), listeners should perceive pauses inserted between prosodic constituents (boundary position) as being more natural than those added within a constituent (non-boundary position). No definition of “more natural” was given, but an example was provided at the beginning of the experiment, followed by a short exposure phase (five items) in which they received feedback for their answers. The participants had to answer using the keyboard, with the key “A” representing the first sentence and key “L” the second sentence. The following five items also consisted of exposure items, but no feedback was given and they were not included in the analysis. The test phase consisted of three blocks (20, 30, 30 items), with the order of the blocks, as well as the order of the items in each block being randomized for every participant.

We also tested non-native (French) speakers, in the original and the synthesized conditions. The original condition would show us how listeners having access only to acoustic information would perform in the task, while the synthesized condition may be seen as a control, to test whether the use of re-synthesis brings any degradation to the acoustic cues marking prosodic boundaries. 22 and 23 participants were tested in sound-proof booths, for each of the two conditions. They had to be monolingual speakers of French, with no knowledge of Japanese. The experimental procedure was identical to that followed for the Japanese participants, the only difference being that the experimenter and the software instructions, as well as the example given at the beginning of the experiment were in French, instead of in Japanese.

### 2.3 Analyses

For each participant, the boundary recognition accuracy was computed using the manual prosodic boundary annotations supplied with the corpus as the reference. The accuracy was defined as the number of correct responses (the boundary position considered to be more natural) divided by the total number of answers given. We then compared the mean accuracy across each condition.

We first performed a linear regression analysis, in order to better understand the effect of the various information sources on the prosodic boundary recognition process. For this, we encoded the five experimental conditions using four dimensions, summarized in Table 1. We then fitted a model with the participant recognition accuracy as dependent variable and the four dimensions as predictors.

The dimension representing the acoustic information that is learned in a language (e.g., its specific tonal markup and the individual weights given to the acoustic cues employed for marking prosodic structure) is called *acousLearned*. The conditions with Japanese participants had this information source enabled (“yes”—Y), while those with French participants did not have it enabled (“no”—N). The dimension *highLevel* encodes whether higher level linguistic information (e.g., lexical, semantic, syntactic) was being exploited in the condition. Only the condition testing Japanese participants hearing the original stimuli had this value set to Y, all the other conditions having it set to N. As this dimension, through the way the higher level linguistic information was removed (re-synthesis or filtering), might have affected also the acoustic signal (e.g., making some acoustic information unavailable or by modifying the original acoustic profile) we introduced a dimension which could help tease apart the effect of this possible loss of information from the overall effect of the *highLevel* dimension. Thus, *nonManip* represents whether the original, non-manipulated, signal was available, the two conditions using the original stimuli having this dimension set to Y and the remaining conditions to N. Last, the dimension

Table 1. The four dimensions used for encoding the experimental conditions and how each condition is being represented by them.

	acousLearned	highLevel	nonManip	phonMore
<i>JP<sub>orig</sub></i>	Y	Y	Y	Y
<i>JP<sub>synt</sub></i>	Y	N	N	Y
<i>JP<sub>filt</sub></i>	Y	N	N	N
<i>FR<sub>orig</sub></i>	N	N	Y	Y
<i>FR<sub>synt</sub></i>	N	N	N	Y

*phonMore* expresses whether more acoustic-phonetic information than the fundamental frequency ( $f_0$ ) of the stimuli is available in the condition. All conditions, but the one in which filtered speech was employed have this value set to Y.

Then, we looked at how similar the judgments of the listeners were across conditions. We assessed this measure by calculating the split-half correlations between conditions as follows: for every listener in the condition, we randomly sampled half of the individual stimuli responses (the same half, for each speaker within a condition) and we computed the accuracy across those stimuli. The same procedure was applied also for the second condition and the Spearman  $\rho$  value was computed on the resulting vectors (containing the per-participant accuracy). We repeated this process 1000 times for each paired comparison and we reported the average  $\rho$  value over all the runs.

### 3. Results

The accuracy of the prosodic boundary recognition, across the five conditions, is presented in Fig. 1. It shows that, as expected, Japanese listeners having access to the original stimuli have a very high recognition performance (mean: 0.932). This performance drops significantly in the conditions employing manipulated stimuli (re-synthesis: 0.658, low-pass filtered: 0.643). These results are numerically similar to those obtained with French listeners in the original stimuli condition (0.648). However, when the re-synthesized stimuli were played to French listeners the recognition rate dropped further, to 0.558. The accuracy obtained in each condition was compared to the chance level (having a value of 0.5) by means of Wilcoxon signed-rank tests. They revealed statistically different performances from chance level, in every condition. Checking the performance difference between conditions, using Wilcoxon rank sum tests, gave the following ranking:  $JP_{orig} > JP_{synth}, JP_{filt}, FR_{orig} > FR_{synt}$ .

We illustrate next the outcome of the linear regression analysis (see Table 2). The intercept represents the mean value of the condition for which all four dimensions are set to Y ( $JP_{orig}$ ). We can see that all dimensions, except for *phonMore* have a significant effect on the recognition of prosodic boundaries. Removing access to the language prosody knowledge encoded by the *acousLearned* dimension reduces the recognition accuracy by 10%. The same for *highLevel*: eliminating information on higher level linguistic information seems to reduce recognition by 18.3%. Since this effect might be confounded with that due to loss of acoustic signal quality (as a consequence of the stimuli manipulation), which also has a significant effect in our analysis (*nonManip*, with an estimate of  $-0.091$ ), the former effect might be, potentially, reduced by up to half. The fact that *phonMore* does not have a significant effect suggests that, for Japanese listeners, the information found in the re-synthesized signal does not bring any complementary information to that found in the filtered signal ( $f_0$ ). Although phoneme duration and speech intensity information is present in the re-synthesized signal, it might be that their quality is not good enough to be reliably employed for boundary recognition. While the language-independent acoustic source of information is not encoded in any of the considered dimensions, we can infer its effect. By comparing the recognition performance in the  $FR_{orig}$  condition (no higher-level linguistic and no language-learned acoustic information used) with the chance level (0.5), one may estimate this effect to be around 0.15.

Then, we looked in more detail at the role of higher-level linguistic information by analyzing whether one of its sub-components, lexical knowledge, has an effect on boundary recognition. As half of the stimuli in each condition had a pause inserted within a word (*syl*) and half between two words (*wrđ*), in the non-boundary condition, we tested the differences between these two positions by means of Wilcoxon signed-rank tests. When a pause was introduced word-internally,

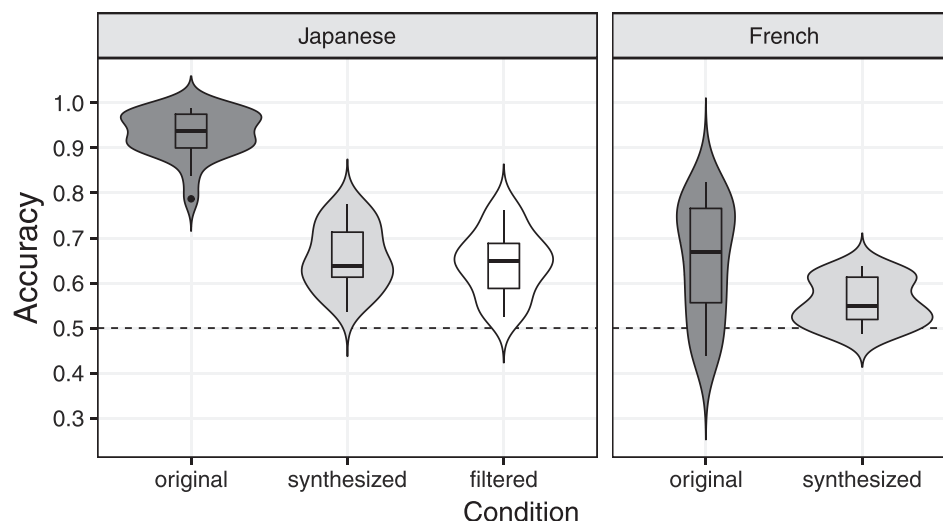


Fig. 1. Boundary recognition accuracy across the five considered conditions. The barplots inside each violin plot illustrate the median and the quartile values. The dashed line represents chance level.

Table 2. Linear regression analysis illustrating the effect of the four dimensions on the boundary recognition accuracy.

Coefficients	Estimate	t-value	p-value
Intercept	0.932	59.67	$< 2e^{-16}$
– acousLearned	–0.100	–4.43	$2.1e^{-5}$
– highLevel	–0.183	–5.71	$8.9e^{-8}$
– nonManip	–0.091	–3.90	$1.6e^{-4}$
– phonMore	–0.014	–0.64	0.524

the mean recognition accuracy in the  $JP_{orig}$  condition was 0.952, while in the case of a pause between two words it stands at 0.911 ( $p = 0.001$ ). This suggests that, indeed, listeners made use of lexical information in their prosodic boundary decision. The differences in the other two conditions involving Japanese speakers went in the opposite direction ( $wrd > syl$ ) and were small:  $JP_{synt}$  ( $wrd$ : 0.661;  $syl$ : 0.654;  $p = 0.586$ ) and  $JP_{filt}$  (0.664; 0.623;  $p = 0.025$ ). Also, for the French participants, the differences in the two conditions followed a similar trend ( $wrd > syl$ ):  $FR_{orig}$  (0.650; 0.647;  $p = 1$ ) and  $FR_{synt}$  (0.578; 0.535;  $p = 0.020$ ). Applying Bonferroni correction to these comparisons ( $p < 0.01$ ), showed that only the differences in the  $JP_{orig}$  condition is statistically significant.

The results of the split-half correlation analysis are shown in Table 3. Besides the pairwise correlations between each two conditions we also report the intra-condition average Spearman  $\rho$  value, for comparison. The intra-condition result may be seen as an upper boundary of the correlation. We observe that most intra-condition correlations have medium positive values ( $\rho > 0.5$ ), implying a rather consistent recognition of boundaries between the participants. The value of the correlation obtained for  $JP_{orig}$  may seem surprising, but it is due to ceiling effects, since many of the split-half correlations in that condition had a value of 1. Comparing the pairwise correlations with the ones obtained intra-condition, we see similar values for the conditions not involving  $JP_{orig}$ , suggesting that both Japanese and French listeners are consistent to each other in their recognition process, when having access to the same type of information. Most of these conditions reach correlations of 0.4 or greater.

#### 4. Discussion and conclusions

We have presented a study investigating the role of different information sources in the perception of prosodic boundaries. Testing native as well as non-native speakers, and using both non-manipulated and manipulated stimuli, we also quantified the effect each of these sources have on the accuracy of boundary recognition. We found that both groups of participants were able to identify prosodic boundaries based only on acoustic cues, significantly better than chance level. Moreover, two acoustic cues shown to mark prosodic boundaries in Japanese, nucleus duration and pitch reset, had a significant role in recognition in all conditions, but the filtered one (see Sec. S5 of the supplementary material). Nevertheless, the addition of other sources of information (higher-level linguistic and other language-specific acoustic information) consistently improved the accuracy, with the maximum performance reached by native speakers listening to the original stimuli (93% accuracy). Our findings are in line with those of previous studies, showing that prosodic boundaries can be discriminated based only on acoustic information [e.g., Carlson et al. (2005) and Kuo (2014)], but also to those indicating that other sources of information play an important role [e.g., Cole et al. (2010), Simon and Christodoulides (2016), and Wakefield et al. (1974)] in recognition. Moreover, our results also suggest that these latter sources provide complementary information to that given by the language-independent acoustic dimension. We were able to quantify not only the effect of these sources of information on the recognition performance, but also the effect of an established de-lexicalization method (re-synthesis).

Due to our experimental design, we could not take into account the role of pauses in the recognition of prosodic boundaries. While the existence of a pause is neither a necessary nor a sufficient condition for the marking of intonation

Table 3. Split-half correlation results (Spearman  $\rho$ ) between the different conditions tested in the study. The colour encodes the strength of the correlation: from light gray (weakest) to black (strongest).

	$JP_{orig}$	$JP_{synt}$	$JP_{filt}$	$FR_{orig}$	$FR_{synt}$
$JP_{orig}$	0.322	0.045	0.176	0.095	0.103
$JP_{synt}$		0.566	0.334	0.372	0.537
$JP_{filt}$			0.616	0.457	0.414
$FR_{orig}$				0.568	0.419
$FR_{synt}$					0.630

phrase boundaries in Japanese (Venditti, 2005), previous work on other languages found pause duration to be the most important acoustic cue used by naïve listeners for postulating a prosodic boundary [e.g., Simon and Christodoulides (2016) for French]. It would be interesting, thus, to further extend this study by taking into account the role of pauses in the acoustic dimension. Moreover, the current study did not consider accentual phrase boundaries. However, it is not clear whether the inclusion of this type of boundary would change the conclusions obtained here. For instance, in the case of the acoustic dimension, although accentual phrases have a distinct tonal profile, which would make them easy to recognize, their boundaries are marked by less extreme values of the other acoustic cues involved. Finally, this work may be expanded to include more detailed analyses, e.g., aiming to further tease apart the role of the individual sources of higher-level linguistic information, by quantifying, for instance, the role of various syntactic categories, similar to Cole *et al.* (2010) and Simon and Christodoulides (2016).

### Acknowledgments

The research of B.L. and E.M. was funded by the European Research Council (Grant No. ERC-2011-AdG-295810 BOOTPHON). It was also supported by the Agence Nationale pour la Recherche (Grant Nos. ANR-10-LABX-0087 IEC, ANR-10-IDEX-0001-02 PSL\*). The work of M.M. and Y.M. was funded by the MEXT Grant-in-Aid for Scientific Research 15H01691 and 20H05010.

### References and links

<sup>1</sup>See supplementary material at <https://www.scitation.org/doi/suppl/10.1121/10.0007150> for further details on the study, the list of employed stimuli, and an acoustic analysis of them.

- Ananthakrishnan, S., and Narayanan, S. (2008). "Automatic prosodic event detection using acoustic, lexical, and syntactic evidence," *IEEE Trans. Audio Speech Lang. Process.* **16**(1), 216–228.
- Carlson, R., Hirschberg, J., and Swerts, M. (2005). "Cues to upcoming Swedish prosodic boundaries: Subjective judgment studies and acoustic correlates," *Speech Commun.* **46**(3-4), 326–333.
- Cole, J., Mo, Y., and Baek, S. (2010). "The role of syntactic structure in guiding prosody perception with ordinary listeners and everyday speech," *Lang. Cogn. Proc.* **25**(7-9), 1141–1177.
- Cutler, A., Dahan, D., and Van Donselaar, W. (1997). "Prosody in the comprehension of spoken language: A literature review," *Lang. Speech* **40**(2), 141–201.
- Johnson, E., Seidl, A., and Tyler, M. (2014). "The edge factor in early word segmentation: Utterance-level prosody enables word form extraction by 6-month-olds," *PLoS One* **9**(1), e83546.
- Kuo, G. (2014). "Processing prosodic boundaries in natural and filtered speech," in *Proceedings of the 7th International Conference on Speech Prosody*, pp. 983–986.
- Ludusan, B., Cristia, A., Martin, A., Mazuka, R., and Dupoux, E. (2016). "Learnability of prosodic boundaries: Is infant-directed speech easier?," *J. Acoust. Soc. Am.* **140**(2), 1239–1250.
- Maekawa, K. (2003). "Corpus of Spontaneous Japanese: Its design and evaluation," in *Proceedings of the ISCA & IEEE Workshop on Spontaneous Speech Processing and Recognition*, p. MM02.
- Pearce, J. W. (2007). "PsychoPy—Psychophysics software in Python," *J. Neurosci. Methods* **162**(1-2), 8–13.
- Simon, A. C., and Christodoulides, G. (2016). "Perception of prosodic boundaries by naïve listeners in French," in *Proceedings of the 8th International Conference on Speech Prosody*, pp. 1158–1162.
- Sloan, R., Akhtar, S. S., Li, B., Shrivastava, R., Gravano, A., and Hirschberg, J. (2019). "Prosody prediction from syntactic, lexical, and word embedding features," in *Proceedings of the 10th ISCA Speech Synthesis Workshop*, pp. 269–274.
- Vaissière, J. (1983). "Language-independent prosodic features," in *Prosody: Models and Measurements*, edited by A. Cutler and D. R. Ladd (Springer, Berlin), pp. 53–66.
- Venditti, J. (2005). "The J\_ToBI model of Japanese intonation," in *Prosodic Typology: The Phonology of Intonation and Phrasing*, edited by S.-A. Jun (Oxford University Press, Oxford), pp. 172–200.
- Wakefield, J. A., Doughtie, E. B., and Yom, B.-H. L. (1974). "The identification of structural components of an unknown language," *J. Psycholing. Res.* **3**(3), 261–269.