CROSSLINGUISTIC VARIATION, PHONETIC VARIABILITY, AND THE FORMATION OF CATEGORIES IN INTONATION

Amalia Arvaniti

University of Kent
a.arvaniti@kent.ac.uk

ABSTRACT

A major issue in intonation research is modelling fine-grained variability while capturing significant generalizations needed to guide typology and abstraction. I argue that this remains an unresolved issue because of the assumed direct and invariant relationship between abstract tonal categories and F0, which is treated as intonation's only exponent. New findings and modelling from my own research programme together with documented typological diversity inform a revised understanding of the relation between abstract intonational structure and phonetic realization. This body of work shows that tonal events are comparable to segments: they are realized by a number of phonetic dimensions that exhibit within-category variability and cross-category overlap, and are in trading relationships with each other. Recognizing the variable realization of tonal events requires that we relax the invariance criterion, and accept that (a) the relationship between intonation and F0 is not straightforward, and (b) intonational meaning is critical for determining intonational categories.

Keywords: intonation, phonetic modelling, typology, phonetic variability

1. INTRODUCTION

The focus of this paper is the representation and phonetic realization of intonation. Intonation refers to the language-specific and systematic modulations of F0 that span entire utterances and have grammatical function(s), such as encoding pragmatic information and marking phrasal boundaries.

A major challenge for models of intonation is the need to simultaneously (i) account for fine-grained phonetic variability in realization and (ii) reach phonological abstractions useful for capturing generalizations about intonation form and meaning. Debates about this relationship are not new or unique to intonation, as the relationship itself applies to all levels of speech analysis. The debate, however, has been particularly prominent in the study of intonation because of additional layers of variability in its phonetic realization and the ways in which they have been handled so far.

First, F0, the primary phonetic exponent of intonation, does not present obvious discontinuities (either at the auditory or acoustic level) that can lead to positing distinct units with some ease. Second, F0 contours treated by speakers and listeners as instances of the same tune can differ substantially from each other, mostly due to context-dependent changes, known as *lawful variability* [5]. These two features of intonation are illustrated in Fig. 1; the F0 contours in black show two instances of uptalk [39]; both could be responses to a prompt like *Name please* and both convey that the speaker is not certain the addressee will recognize their name. However, the final rise is timed differently, spanning the last two syllables in the longer utterance, but just a part of the only vowel in the shorter one. The blue line shows the pitch contour of a different token of *Leonard Millersville* produced by the same speaker with comparable duration to the token with the black contour. When considering the differences between these three pitch tracks, three questions arise. First, how do we capture the commonalities that lead speakers of some English varieties to interpret these contours as instances of the same tune? Second, how do speakers know when to use this tune, and how do they learn how to implement it over utterances of different length and structure? Third, as linguists, how do we separate (inevitable) random noise, like that between the blue and black contours in Fig. 1, from other types of variability that we may need to account for, such as dialect, style and politeness effects [4, 14, 15, 20], and how do we separate these from inter-speaker variability [14]? In short, the challenge is how to make sense of all these types of variability, how to abstract generalities beyond them, and how to disentangle phonological intonational contrasts as encoded in F0, from other influences on F0.

Figure 1: Waveforms and F0 contours of uptalk used with utterances differing in duration; the blue and black lines represent the contours of two tokens of *Leonard Millersville* produced by the same speaker.

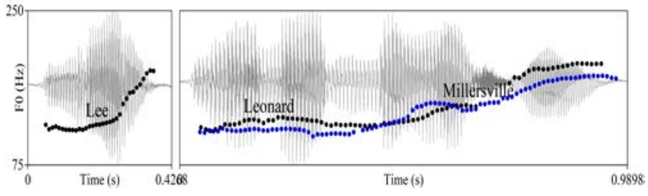

1.2. Treatments of phonetic variability and abstraction

The challenge of determining tonal constituents and structure while handling variability has met with a number of approaches. Most focus on one of the challenges only. Thus, several early models developed representations that are by and large idealized versions of pitch contours, and (with the exception of IPO [38]), do not address variability at all [10, 19, 30, 33]. This is illustrated in Fig. 2 using the British School notation [30].

Figure 2. An idealized pitch contour using the British School notation; from [30, p. 29].

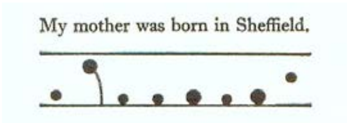

Some models have instead focused on capturing variability [13, 34]. For example, Xu and colleagues used root mean square error (RMSE) to measure the distance between natural F0 curves and curves synthesized following the principles of a number of models [34]. In this approach, the goal is to be as close as possible to the original F0 contour and capture every microprosodic effect. As argued in [6], by focusing exclusively on granularity, these types of models fail to capture significant generalizations. At the same time, they fail to separate lawful variation from random variability.

In contrast to the above models, the autosegmental metrical model of intonational phonology (henceforth AM [24]) appears at first glance to address both challenges. AM is a phonological model in which intonation is represented as a string of H (high) and L (low) tones phonologically associated with the prosodic tree, specifically with either prosodic heads (typically stressed syllables) or phrasal boundaries. The model provides a mechanism for connecting the abstract tonal string to phonetics: L and H tones are phonetically realized as *tonal targets*, local minima and maxima defined by two dimensions, their *scaling*, i.e. their pitch (or F0) height, and their *alignment*, i.e. their synchronization with the segmental string. This synchronization reflects the phonological association of the tones with prosodic structure: e.g., H* is a pitch accent, i.e. a tone that associates with a prosodic head; phonetically this H* is realized as high F0 on a stressed syllable (or close to it).

Although AM would appear to provide a solution to the main challenges noted above [6], a number of assumptions, practices, and diagnostic criteria have conspired to turn variability into a problem.

First, although AM is a phonological model, its representations are often treated as akin to phonetic *transcriptions*. This understanding in effect treats the tonal string as a reconstruction of F0 (like other idealizations), not a symbolic representation. It has thus led to an assumption that AM representations should be faithful reflections of F0, thereby sustaining the confound between F0 and intonation (see [24] for a discussion on how this assumption may have been strengthened by the popularity of ToBI systems). The use of L and H to represent tones reinforces this view in a way that does not apply to symbols for segments. The relation between, say, [a] and its phonetic value is recognized to be one of convention, and [a] can be used to represent a number of related vowel qualities. For L and H, however, there is an expectation (which is, strictly-speaking, unwarranted by AM), that they transparently represent actual low and high F0 points [24].

An additional issue is the notion of *segmental anchoring* [6], which postulates that tonal targets reflecting underlying tones align with the segmental string in a stable manner. This finding of [6] has led to stable alignment being used as a diagnostic of phonological status. For instance, [25] use the stable scaling and alignment of an F0 minimum to argue that the English pitch accent represented as H* in [31] includes a low tonal target, and thus that this accent's phonological representation should include a L tone.

More recently, the preoccupation with variability in the face of expected invariance has led to the development of an International Prosodic Alphabet (IPrA), in which the precise location of F0 minima and maxima is a primary analysis criterion [21]. Fig. 3 illustrates how two rises with slightly different peak alignment are represented in [21]. The assumption is that L+H* and L+<H* could be contrastive in a system. For such precision to work in practice, however, tonal targets must be invariant.

Figure 3: IPrA representations of the alignment contrast between L+H* and L+<H* accents; black lines represent idealized F0; blue rectangles represent the stressed syllable; after [21].

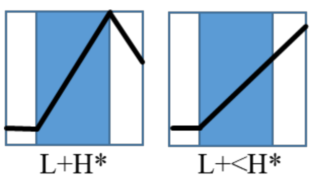

Given many standard practices in intonation research (see e.g. [22, 23]), the assumption of invariance evinced in [21] may seem well-founded. Many studies rely on scripted data elicited from dialectally homogeneous groups of educated speakers who are at ease with reading from a script in a consistent manner; thus the data of many existing studies are relatively uniform [2]. Further, statistical analysis focuses on the presence of significant differences between putative categories, since such differences are the main criterion for determining tonal status. These practices create the impression that tonal categories are (or should be) phonetically distinct and invariant. They lead to an expectation that distributions for some phonetic feature, such as early vs. late peak alignment said to distinguish L+H* and L+<H* in Fig. 3, are akin to the distributions in Fig. 4a: narrow and completely distinct from each other. On the other hand, realistic distributions of the same feature (peak alignment), like those in Fig. 4b, are considered problematic.

Figure 4. (a) Idealized distributions of early peak (EP) vs. late peak (LP) alignment; (b) schematic visualization of realistic distributions; (c) peak alignment as a continuum from early to late peak.

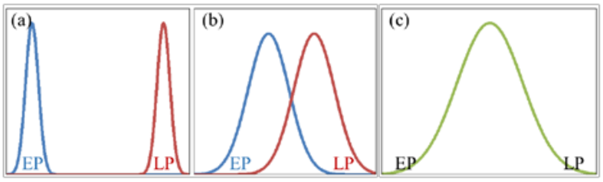

In short, the approaches to intonation so far have either ignored variability altogether, or have made it into a core issue: it is either the main focus of modelling or a problem to be solved. In what follows I present results from recent work which show why this approach should be revised. These results support a different approach to intonation and advocate in favour of greater reliance on meaning.

2. PHONETIC VARIABILITY AND CATEGORY OVERLAP

The extent of variability in intonation is glimpsed by means of Functional Principal Components Analysis (FPCA, [17]). FPCA captures essential modes of variation in curves and returns them in functional form, as Principal Components (PCs). Each F0 curve in the analysed sample receives a coefficient for each PC, showing the extent to which this PC contributes to the shape of this curve. In [26] we used FPCA to analyse data from thirteen speakers of Greek who produced short utterances in dialogues leading to the use of one of three accents, H*, L+H* and H*+L, depending on pragmatics [3]. The accents show variability due to tonal crowding: they all appear in one of the last three syllables of an utterance followed by L-L% edge tones.

FPCA of 844 F0 curves that constituted this corpus showed that the first two PCs capture 88.7% of the variability in the curves. The shape of the PCs is illustrated in Fig. 5: the black line represents the average curve in this body of data; the curves with + and – symbols represent the changes to the average curve when the coefficient of the depicted PC is +1 or -1 standard deviation respectively. F0 curves are composites of PC1 and PC2 (and additional PCs to a much lesser extent), with the contribution of each PC to the shape of the F0 curve being determined by its coefficient. In their aggregate the PC coefficients (or *scores*) differ significantly by accent (for the statistical analysis see [26]). What is of interest here is the overlap between accents in terms of these scores. This is illustrated in the density plots of PC1 and PC2 in Fig. 6, which shows substantial overlap in PC scores for the three accents.

Such overlap should not be surprising. It is normal for all types of phonetic categories. It applies even to categories that are critical for encoding phonological contrasts, such as short- and long-lag VOT which phonetically distinguish phonologically voiced and voiceless stops in English [28]. The PC density plots also show that some categories may have more stable realization than others; this is the case here with the H* accent, relative to L+H* and H*+L [cf. 37].

Further, the data show that scaling and alignment are not independent of one another: as the curves in Fig. 5 show, peak location and scaling covary; e.g. a higher PC1 score leads to an earlier and higher peak relative to a low PC1 score; the peak of the positive PC2 curve is both higher and delayed relative to the peak of the negative curve. This suggests that it is not possible to neatly distinguish alignment from scaling and consider them independent of each other.

Figure 5. PC1 and PC2 curves modelling 844 F0 curves evenly distributed among the H*, L+H* and H*+L accents of Greek (for details, see text and [26]).

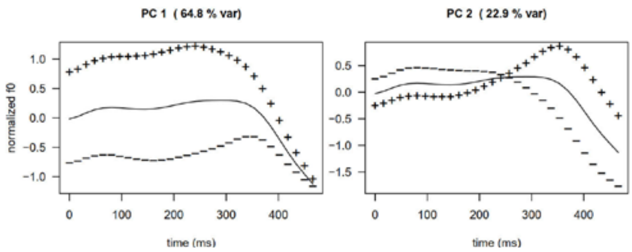

Figure 6: Density plots of PC1 scores (a) and PC2 scores (b), separately for H*, L+H*, and H*+L Greek accents.

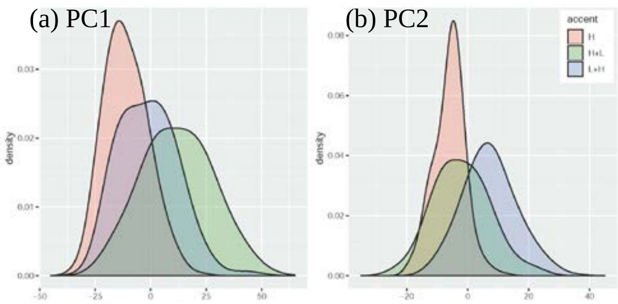

Such variability is not unique to these accents; similar results are presented, e.g., for Greek wh-questions [15, 16], despite the fact that the investigated differences in pitch contours are pragmatically robust [9]. More generally, variability has been reported for a number of languages. For example, loose tonal alignment is a characteristic of several languages, including Dalabon [11], Mawng [12], and Ambonese Malay [27]. These findings indicate that strict alignment is not necessary, and that a looser association between tonal events and structural positions may be the norm in some systems. This could be due to other prosodic properties, such as a lack of lexical stress, as in Ambonese Malay [27], or to word order and morphological structure [11, 12].

3. BEYOND ALIGNMENT

As the above findings indicate, overlap between tonal categories is inevitable, and stable alignment may not be an option at all for some categories or linguistic systems. If so, it is worth considering alternative measures, criteria and practices in the analysis of intonation.

3.1. Multiple cues and cue trading

One of the reasons for the focus on invariance is that F0 is considered the only exponent of intonation and thus as essential for determining the identity of a tonal event. However, many studies show that tonal events influence the realization of segments they co-occur with. This suggests that tonal events are cued by multiple parameters, not just F0 [8, 29]. Cue redundancy and cue-trading are not new concepts when dealing with segments [35], though they are only beginning to be discussed in intonation [14].

Such cues can be specific to particular tonal events. For instance, the PC scores of the Greek pitch accents discussed above have been statistically modelled to include the duration of the analysis window. All accents showed a negative correlation between duration and PC1, suggesting a trade-off between pitch height and duration to encode accentuation. However, for L+H*, the results also showed a negative correlation between PC2 and duration; in other words, for this accent there was a trade-off between getting the "scooped" shape right or using duration to encode contrastiveness (cf. [7] on the role of duration in realizing this accent). A similar connection is reported in [2] for the L+H* of Greek Thrace Romani. Likewise, [8] examined two tunes used in Polish for routine and urgent calls and report than the latter is associated with higher RMS amplitude of the accented syllable. Similar results are reported in [15] for the wh-question tunes of Greek. In [15] it is shown that although accentual peak alignment (early or late) and boundary tone scaling (!H% or L%) are both strong tune predictors, so is the duration of the final vowel.

Additional cues need not be local to the syllable(s) synchronized with a tonal event. The Greek accent data were also modelled to include the presence of a preceding accent and its effect of PCs [26]. The results show that the presence of a preceding accent does affect F0 scaling, with effects being specific to each accent and consistent with its representation [26]. These changes, which affect the F0 of *unaccented* syllables, could be cues that help listeners predict the identity of upcoming tonal events. Although the role of these additional cues is as yet unclear, it is evident that they are worthy of further investigation.

3.2. Optional tonal targets

An additional criterion for phonological status is proposed by [8]. The authors examined routine and urgent calls in Polish and showed that the rise found in both is treated differently under tonal crowding. As shown in Fig. 7, when these tunes are realized on a monosyllable, the rise to a peak in the routine tune is retained, while that of the urgent tune is truncated. This suggests that the rise is essential for the identity of the accent only in the former tune. This is reflected in the representations proposed in [8], namely LH* for the routine and H* for the urgent call accent.

Figure 7. Waveforms and F0 of *Piotr* with the routine (left) and urgent call (right) tune of Polish; after [8].

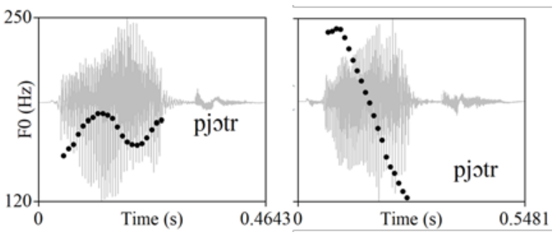

3.3. Analytical decisions

A practical question that arises is how to deal with the variability documented above when analysing a new intonation system. In [2] I discuss a number of sources of variation and argue that they should be considered before observed changes are deemed to reflect a phonological difference. In short, in [2] I advocate for analytical simplicity unless evidence suggests otherwise.

This principle is put to practice in the analysis of Greek Thrace Romani intonation [2]. The corpus is one of spontaneous speech from a non-standardized linguistic variety spoken primarily by speakers with little schooling who are trilingual in Romani, Turkish and Greek. Under these conditions, typical tests of invariance fail. This is illustrated in Fig. 8, which shows several instances of a H* accent in an utterance from a semi-spontaneous task. No two instances of

H* are similar phonetically, with differences relating to segmental context and the vicinity of other accents. The solution advocated in [2] is to rely on the pragmatic role of the accents: the utterance in Fig. 8 is all new, so the accents are treated as instances of the same tonal category. To ensure parsimony, these accents are represented as H*, as the presence of a preceding low target is optional. What the H* tokens have in common in terms of phonetics is that the accented syllable has high F0 in all of them.

Figure 8. Waveform and F0 contour of “on Abdul there are three apples” showing multiple instances of H* accents in Greek Thrace Romani; after [2].

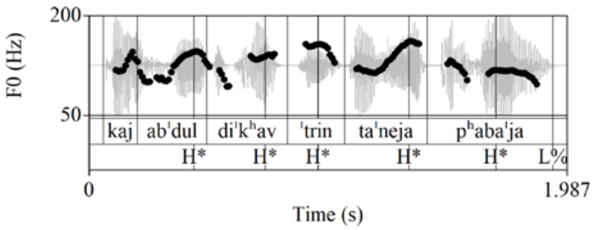

4. CONTINUA, CATEGORY OVERLAP AND THE ROLE OF MEANING

This understanding of phonetic realization requires systematic use of intonational meaning as a criterion during analysis. Although formal models of intonational meaning are available [32, 36], they have not been extensively tested or systematically applied to languages other than English. In addition, many descriptions of intonational meaning blur the boundaries between linguistic meaning and paralinguistic information. A case in point is the debate as to whether H* and L+H* are two distinct accents in English or present the edges of a continuum (and thus they are instances of the same category). In [32], the analysis that posits two categories, H* and L+H*, the meaning of the accents is formalized in terms of their contribution to the common ground, i.e. within a recognized pragmatic framework. In alternatives such as [24], the accents form a continuum that stretches from least to most emphatic, with earlier peak alignment (and lower scaling) at one end and late peak alignment (and higher scaling) at the other; in short, this leads to a distribution like that in Fig. 4c, along the scaling and alignment dimensions. As argued in [1], the analysis in [32] relates to *intonation*, while that in [24] is paralinguistic: both the realization and the interpretation are gradient. The former analysis allows for within-category variability (as in Fig. 4b), while the latter can only reflect genuine gradience. The possible presence of such gradience and its relation to intonational contrasts and within-category variability is a topic worth examining further.

5. TINT: TOWARDS A NEW APPROACH

Below I propose a new approach to intonation research based on the above findings. It is a set of a principles, stemming from a main postulation, which can be condensed into *tame intonation (TINT):* intonation is comparable to other subsystems of phonology, not Bolinger’s “half-tamed savage” [10]. The “untamed” side of intonation reflects paralinguistic, non-structural effects on F0, which, should not be confused with intonation proper. It follows that the study of intonation should rely on the criteria and assumptions used for the study of segmentals. A number of corollaries stem from TINT.

Realization

- Intonational categories are likely to show phonetic variability the extent and sources of which can only be empirically determined.
- Although F0 is the main exponent of intonation, it is unlikely to be the only one.
- Not all F0 variations reflect phonological structure; empirical testing is required to determine which do.

Representation and analysis

- Phonological representations should be used to capture contrastiveness in a system, not variability.
- Criteria for positing contrasts should include phonetic evidence, system-internal considerations [18], and the role of meaning.
- Invariance on some phonetic dimension should not be considered essential or used as the main criterion for establishing phonological categories.
- Paralinguistic notions (such as emphasis) should not be used in lieu of formal pragmatic definitions of meaning distinctions.

Processing

- The identification and interpretation of tonal categories during processing relies on (i) local phonetic detail and cue-weighting, (ii) distal context, such as the realization of other tonal events in the utterance, and (iii) top-down information, including pragmatic context.

6. CONCLUSION

In conclusion, this body of work invites a re-think of standard practices and assumptions in the study of intonation that go beyond AM. It shows that tonal events are comparable to segments: they are realized by several phonetic dimensions that are in trading relationships with each other and exhibit within-category variability and cross-category overlap. Recognizing the variable realization of tonal events requires that we relax the invariance criterion, and accept that the relationship between intonation and F0 is not straightforward, and intonational meaning is critical for determining intonational categories.

7. REFERENCES

[1] Arvaniti, A. 2007. On the relationship between phonology and phonetics (or why phonetics is not phonology). *Proceedings of ICPhS XVI*, pp. 19–24.

[2] Arvaniti, A. 2016. Analytical decisions in intonation research and the role of representations: Lessons from Romani. *Laboratory Phonology*: *Journal of the Association for Laboratory Phonology* 7(2), 6. DOI: http://dx.doi.org/10.5334/labphon.14

[3] Arvaniti, A. & M. Baltazani. 2005. Intonational analysis and prosodic annotation of Greek spoken corpora. In: S. Jun (ed). Oxford: OUP, 84–117.

[4] Arvaniti, A. & G. Garding. 2007. Dialectal variation in the rising accents of American English. In: Cole, J., Hualde, J. I. (eds), *Papers in Laboratory Phonology 9*. Mouton de Gruyter, 547–576.

[5] Arvaniti, A. & D. R. Ladd. 2009. Greek wh-questions and the phonology of intonation. *Phonology* 26, 43-74.

[6] Arvaniti, A., D. R. Ladd & I. Mennen. 1998. Stability of tonal alignment: The case of Greek prenuclear accents. *Journal of Phonetics* 26, 3–25.

[7] Arvaniti, A., D. R. Ladd & I. Mennen. 2006. Tonal association and tonal alignment: evidence from Greek polar questions and contrastive statements. *Language and Speech* 49: 421–450.

[8] Arvaniti, A., M. Żygis & M. Jaskuła. 2017. The phonetics and phonology of the Polish calling melodies. *Phonetica* 73(3-4), 338–361.

[9] Baltazani, M., S. Gryllia & A. Arvaniti. 2019. The intonation and pragmatics of Greek wh-questions. *Language and Speech*. DOI: 10.1177/0023830918823236.

[10] Bolinger, D. L. 1986. *Intonation and Its Parts: Melody in Spoken English*. London: Arnold.

[11] Fletcher, J. 2015. Intonation and prosody in Dalabon. In: Jun, S. (ed), Oxford: OUP, 252–272.

[12] Fletcher, J., Stoakes, H., Singer, R., Loakes, D. 2016. Intonational correlates of subject and object realization in Mawng (Australian). *Proc. Speech Prosody 2016.*

[13] Fujisaki, H. 1983. Dynamic characteristics of voice fundamental frequency in speech and singing. In: MacNeilage, P. F. (ed), *The production of speech*. Heidelberg: Springer-Verlag, 39–55.

[14] Grice, M., Ritter, S., Niemann, H., Roettger, T. B. 2017. Integrating the discreteness and continuity of intonational categories. *JPhon* 64: 90–107.

[15] Gryllia, S., Baltazani, M., Arvaniti, A. 2018*a*. The role of pragmatics and politeness in explaining prosodic variability. *Proc. Speech Prosody 2018.*

[16] Gryllia, S., Baltazani, M., Arvaniti, A. 2018*b*. Variability and multiplicity of cues in the realization of intonation. LabPhon 16, Lisbon, 19-22 June 2018.

[17] Gubian, M., Torreira, F., Boves, L. 2015. Using functional data analysis for investigating multidimensional dynamic phonetic contrasts. *JPhon* 49, 16–40.

[18] Gussenhoven, C. 2016. Analysis of intonation: The case of MAE_ToBI. *Lab Phon 7*(1), 10. DOI: http://doi.org/10.5334/labphon.30

[19] Hirst, D., Di Cristo, A., Espesser, R. 2000. Levels of representation and levels of analysis for the description of intonation systems. In: Horne, M. (ed), *Prosody: Theory and Experiment, Studies Presented to Gösta Bruce*. Dordrecht: Kluwer, 51–87.

[20] Henriksen, N. 2013. Style, prosodic variation, and the social meaning of intonation. *JIPA* 43(2), 153–193.

[21] Hualde, J. I., Prieto, P. 2016. Towards an International Prosodic Alphabet (IPrA). *Lab. Phon.* 7(1), 5. DOI: http://doi.org/10.5334/labphon.11

[22] Jun, S.-A. (ed). 2005. *Prosodic Typology: The Phonology of Intonation and Phrasing.* Oxford: Oxford University Press.

[23] Jun, S.-A. (ed). 2014. *Prosodic Typology II: The Phonology of Intonation and Phrasing.* Oxford: Oxford University Press.

[24] Ladd, D. R. 2008. *Intonational Phonology*. CUP.

[25] Ladd, D. R. & A. Schepman. 2003. ‘‘Sagging transitions’’ between high pitch accents in English: Experimental evidence. *JPhon* 31, 81–112.

[26] Lohfink, G., Katsika, Arvaniti, A. 2019. Variability and category overlap in the realization of intonation. *Proceedings of ICPhS 2019*.

[27] Maskikit-Essed, R., Gussenhoven, C. 2016. No stress, no pitch accent, no prosodic focus: The case of Ambonese Malay. *Phonology* 33, 353–389.

[28] Nakai, S., Scobbie, J. M. 2016. The VOT category boundary in word-initial stops: Counter-evidence against rate normalization in English spontaneous speech. *Laboratory Phonology* 7(1), 13. DOI: http://doi.org/10.5334/labphon.49.

[29] Niebuhr, O. 2012. At the edge of intonation–The interplay of utterance-final F0 movements and voiceless fricative sounds in German. *Phonetica* 69, 7–21.

[30] O’ Connor, J. D., Arnold, G. F. 1973. *Intonation of Colloquial English*. London: Longman.

[31] Pierrehumbert, J. B. 1980. *The Phonology and Phonetics of English Intonation*. MIT dissertation.

[32] Pierrehumbert, J. B., Hirschberg, J. 1990. The meaning of intonational contours in the interpretation of discourse. In: Cohen, P. R., Morgan, J. L., Pollack, M. E. (eds), *Intentions in Communication*. The MIT Press, 271–311.

[33] Pike, K. L. 1945. *The Intonation of American English*. Ann Arbor, MI: U. of Michigan Press.

[34] Prom-on, S., Xu, Y., Gu, W., Arvaniti, A., Nam, H. Whalen, D. H. 2016. The Common Prosody Platform, (CPP) – where theories of prosody can be directly compared. *Proc. Speech Prosody 2016*.

[35] Repp, B. H. 1982. Phonetic trading relations and context effects: New experimental evidence for a speech mode of perception. *Psychol Bull* 92: 81–110.

[36] Steedman, M. 2014. The surface-compositional semantics of English intonation. *Lg* 90(1), 2–57.

[37] Schweitzer, K., et al. 2015. Exploring the relationship between intonation and the lexicon: Evidence for lexicalised storage of intonation. *Speech Commun* 66, 65–81.

[38] ’t Hart, J., Collier, R., Cohen, A. 1990. *A Perceptual Study of Intonation*. Cambridge: C.U.P.

[39] Warren, P. (2016). *Uptalk: The Phenomenon of Rising Intonation*. Cambridge: Cambridge University Press.