Encoding and decoding of meaning through structured variability in intonational speech prosody

Short title: Structured variability in speech prosody

Andrés Buxó-Lugo*

Chigusa Kurumada

Department of Brain and Cognitive Sciences, University of Rochester

Rochester, NY, 14627 USA

*Please address correspondence to

Andrés Buxó-Lugo

Department of Brain and Cognitive Sciences,

University of Rochester,

303B Meliora Hall,

Rochester, NY 14627

Email: abuxolug@ur.rochester.edu

Abstract

Speech prosody, the rhythm and intonation in particular, plays an important role in communication of meaning. Rising vs. falling intonation contours signaling the speaker's intended communicative meanings (i.e., asking a question vs. making a statement) has been widely recognized as a primary example of such prosody. However, what appears to be a straightforward mapping between acoustic features of prosody and hypothesized meanings in fact presents a challenge to the human perceptual and computational mechanisms. Perceptible features of prosody vary across contexts (e.g., talkers), creating ubiquitous ambiguity in the mapping. Here, we first characterize the structured nature of the variability in intonational speech prosody used to signal a question vs. a statement in English. We then demonstrate that the listener can learn to adapt their expectations about the prosody-meaning mapping according to an inferred underlying structure of the environmental input. We argue that the rich and dynamic representations of the prosodic input allow listeners to infer the mapping that is most likely given characteristics of a current context.

Keywords: Prosody, meaning, intonation, adaptation, language production, language comprehension, variability

1. Introduction

A fundamental aspect of human language is its ability to communicate meaning. Among various devices used to convey meaning, speech prosody critically supports communication of linguistic as well as social meaning, such as the talker's intentions and emotions. Across languages, various temporal and rhythmic features (e.g., stress, accent) distinguish lexical identities (e.g., Cutler, 2015; Dahan, 2015; Ladd, 2008 for reviews), and intonation patterns can carry information about grammatical structure (e.g., Cooper & Paccia-Cooper, 1990; Ferreira, 1993; Frazier, Carlson, & Clifton, 2006; Price, Ostendorf, Shattuck-Hufnagel, & Fong, 1991; Schafer, Speer, Warren, & White, 2000, among others). Furthermore, prosodic junctures and emphasis signal information structure and statuses of discourse referents (e.g., Arnold & Watson, 2015; Bolinger, 1989; Bolinger, 1961; Braun & Ladd, 2003; Büring, 2003; Calhoun, 2010; Swerts & Geluykens, 1994; Venditti & Hirschberg, 2003) as well as their interpretations with respect to information shared among interlocutors (Dennison & Schafer, 2010; Hirschberg, 2002; Ito & Speer, 2008; Pierrehumbert & Hirschberg, 1990; Tomlinson, Gotzner, & Bott, 2017). For instance, the locations and types of pitch accents signal information status of referents in English (e.g., “Fred ate the BEANS” (in response to “What did FRED eat?”) vs. “FRED ate the beans” (in response to “WHO ate the beans?”) (Büring, 2003). Prosody thus constitutes a unique interface between perceptual properties of speech and grouping of linguistic representations bearing relative prominence in context (Cole, 2015; Pierrehumbert & Hirschberg, 1990; Wagner & Watson, 2010).

Despite a long history of research, empirical investigations of prosody have remained one of the most impenetrable areas of speech science. This is primarily due to

the daunting complexity and variability in how the prosodic signal is structured and used (Brown, Newcomb, & Kahrl, 1995; Cole & Shattuck-Hufnagel, 2016; Dilley, Arjmandi, Ireland, Heffner, & Pitt, 2016; Kurumada, Brown, & Tanenhaus, 2017; Lennes, Stevanovic, Aalto, & Palo, 2016; Zipp & Dellwo, 2011). Actual realizations of a prosodic event (e.g., a pitch escalation) vary along many acoustic dimensions (e.g., fundamental frequency, duration, and intensity), and the same event can be attributable to multiple possible meanings and interpretations (e.g., conveying a lexical stress, informational emphasis, talker intentions such as surprise and incredulity). Prosody also varies according to a talker's socio-indexical features such as talker identities, accents (Arvaniti & Garding, 2007; Clopper & Smiljanic, 2011; Grabe, Kochanski, & Coleman, 2005; Ladd, 2008), age (Horton, Spieler, & Shriberg, 2010; Patel & Brayton, 2009; Patel & Grigos, 2006), gender (Clopper & Smiljanic, 2011; Fitzsimons, Sheahan, & Staunton, 2001), and social groups (Green, 2002; Holliday, 2019; Podesva & Callier, 2015; Warren, 2016). Some of the variation is tied to characteristics of the talker (e.g., a mean fundamental frequency (f_0) correlating with the vocal fold length (Titze, 1989)) while other variation is deployed as a means for stylistic effects (Gryllia, Baltazani., & Arvaniti, 2018). Furthermore, even within each speaker, prosodic production varies across speaking conditions, styles, and formality (Hawkins, 2003; Henriksen, 2013), as well as intended audience (e.g., child- vs. adult-directed speech, e.g., Dilley, Millett, McAuley, & Bergeson (2014); Fernald & Kuhl, (1987); Foulkes, Docherty, & Watt, (2005)).

One key question stemming from the variability of prosody is the gradient nature of acoustic correlates that support prosody-meaning mappings (Dilley, 2010; Ladd &

Morton, 1997; Shattuck-Hufnagel, Dilley, Veilleux, Brugos, & Speer, 2004). An early and influential proposal suggested that, to extract meaning from prosody, listeners map continuous, acoustic cues onto hierarchically nested, discrete, prosodic categories (e.g., word, foot, intonation phrases) (e.g. Beckman & Ayers Elam, 1997; Beckman, Ayers, & Ayers Elam, 1994; Bruce, 1977; Goldsmith, 1976; Pierrehumbert, 1980). For instance, in English, f_0 , duration, and intensity values combine to support syllable-level (e.g., lexical stress) to phrasal-level (e.g., phrasal tone) categories, which serve as operators to link acoustic features of speech to meanings relevant to discourse (Pierrehumbert & Hirschberg, 1990). As discussed above, however, acoustic realizations of a prosodic category can be variable due to random factors, such as errors in production and perception, as well as due to systematic factors such as differences between linguistic contexts, utterance length, talkers, dialects/accents, and speech conditions. In other words, the same physical cues can be associated with a particular prosodic (phonological) category with varying degrees of (un)certainly across contexts. This prevalent variability has been a major roadblock for building a phonological theory of speech prosody (Breen, Dilley, Kraemer, & Gibson, 2012; Krahmer & Swerts, 2001; Pitrelli, Beckman, & Hirschberg, 1994; Syrdal & McGory, 2000; Yoon, Chavarria, Cole, & Hasegawa-Johnson, 2004).

Past research has begun to answer this question through two, not mutually exclusive, mechanisms. One possibility is that listeners normalize acoustic cues according to a contextually defined baseline (Dilley & Pitt, 2010; Francis, Ciocca, Wong, Leung, & Chu, 2006; Leather, 1983; Moore & Jongman, 1997; Newman & Sawusch, 1996; Pisoni, 1977; Zipp & Dellwo, 2011). Both behavioral and neuro-imaging results

support that perception of pitch contours is conditioned on an expected baseline for a speaker and a speaker group (e.g., males vs. females or adults vs. children (Bishop & Keating, 2012; Tang, Hamilton, & Chang, 2017)). Such normalization can occur rapidly, even within a matter of 20 to 50 ms (Lee, 2009). Similarly, durational information can be normalized according to a contextual speech rate (Baese-Berk et al., 2014; Diehl, Souther, & Convis, 1980; Newman & Sawusch, 2009; Sawusch & Newman, 2000). Information received at an earlier part of the sentence can dynamically influence interpretations of relative duration of syllables downstream in the same utterance (the so-called “distal prosody effect”) (Baese-Berk et al., 2014; Brown, Salverda, Dilley, & Tanenhaus, 2015; Dilley & Pitt, 2010; Morrill, Dilley, McAuley, & Pitt, 2014; Reinisch & Maximilian, 2015).

This manner of reasoning is undergirded by the past finding that normalization of speech sounds is a basic human perceptual mechanism that operates at low-level auditory processes (Adank, Nuttall, Banks, & Kennedy-Higgins, 2015; Adank, Smits, & van Hout, 2004; Flynn & Foulkes, 2011; McMurray & Jongman, 2012). Recognition of a sound is considered to include automatic transformation of acoustic values according to multiple instances of the same category (e.g., vowel) as the current input (i.e., extrinsic normalization) or other features perceived concurrently (intrinsic normalization e.g., standardizations of F1 against f_0 in vowel perception). Often implicit in this view of acoustic normalization is the assumption that there is an underlying, invariant, mapping between acoustic cues and abstract linguistic categories beneath the surface variations (e.g. Johnson, 2005; Lobanov, 1971; Monahan & Idsardi, 2010). Variations need to be

stripped away through normalization so that the listener can arrive at abstract representations universal across talkers and contexts.

An additional way of coping with the variability in the cue-category mapping is to store perceived exemplars (or parametric summaries thereof) and learn characteristics of prosodic production specific to a context (e.g., Hawkins, 2003; Pierrehumbert, 2001; Schweitzer, 2012). Here, we take a *talker identity* as an example of a context for prosodic interpretations. Upon encountering a talker who frequently uses an upward intonation for a declarative sentence (e.g., “uptalk”; Warren, 2016), listeners can accommodate the idiosyncrasy and process subsequent utterances according to the previously heard exemplars. Similarly, it has been shown that accumulated exposure to a specific talker allows listeners to accurately predict a terminal intonation contour at an earlier part of an utterance (Saindon, Trehub, Schellenberg, & Van Lieshout, 2017). An emergent body of work has also suggested that listeners can learn distributional statistics of acoustic features associated with a prosodic unit (e.g., intonation contour). Kurumada, Brown, & Tanenhaus (2017) exposed native English talkers to intonation contours sampled from a continuum ranging between “It looks like a ZEBRA” (pronounced with a nuclear accenting on the final noun and a falling boundary tone) and “it LOOKS like a zebra...” (pronounced with a nuclear accenting on the verb and a rising boundary tone). These two contours support two contrasting interpretations (i.e., (It looks like a zebra and) *it is one* vs. *it is something else*). Listeners who received distinct sampling patterns of the acoustic cue values for these two contours subsequently provided diverging interpretations for a novel token mid-way between the two ends of the continuum (cf. perceptual learning paradigm in speech perception (Norris, McQueen, & Cutler, 2003; Samuel & Kraljic,

2009)). This and related results suggest that listeners are able to learn characteristics of the talker and/or a context to adapt their interpretations accordingly (Roettger & Franke, under review).

Because of distinct computational assumptions and goals of applications, normalization and learning (adaptation) of prosody have so far been examined separately. This separation leaves open a critical question regarding how listeners navigate the variability prevalent in the human speech sounds to infer meanings encoded in prosody. To fill this gap, the current study first examines whether, and to what extent, the acoustic features of speech can be reliably mapped onto prosodic categories after a normalization method is applied. If there is a transformation function that supports an invariant mapping between acoustic cues and prosodic categories, that would motivate a broader question of whether such a function is hardwired in our auditory system or can be learned from the data. If, alternatively, the prosodic input is structured in a manner so that no unique function can be posited, or if distinct functions need to be applied across contexts, that would prompt us to explore how listeners can learn a cue-category mapping for each new context (e.g., a novel talker) (Weatherholtz & Jaeger, 2016). An overarching goal of the current study is thus to uncover the nature of prosodic variability to better understand challenges the human language processing system faces in the interpretation of meaning encoded through the channel.

As an interesting test case, we chose to explore English listeners' interpretations of utterance final intonation contours, often associated with two contrasting communicative meanings, or speech acts (i.e., asking a question vs. making a statement). This is a well-recognized and thoroughly investigated contrast, readily accessible to

native listeners (Couper-Kuhlen, 1984; Doherty, West, Dilley, Shattuck-Hufnagel, & Caplan, 2004; Hedberg, Sosa, & Görgülü, 2017; Ladd, 1996; Studdert-Kennedy & Hadding, 1973). Experimental evidence suggests that the meanings are derived in an incremental, and sometimes predictive, manner as an utterance unfolds over time (Bögels & Torreira, 2015; Heeren, Bibyk, Gunlogson, & Tanenhaus, 2015; Saindon et al., 2017). Despite the seemingly straightforward nature of the mapping between acoustic features and prosodic categories (e.g., rising vs. falling that support contrasting meanings), exact mechanisms supporting the mapping remain to be debated. Most important, the distinction between rise vs. fall is inherently a relative notion (paraphrased to “higher” or “lower” with respect to some baseline). What counts as rise vs. fall critically depends not only on proximal spectral and durational cues but also on those that are distal to a terminal contour (e.g., F0 of a stressed syllable, Studdert-Kennedy & Hadding, 1973). Furthermore, acoustic realizations of rising vs. falling contours (as indicative of a question vs. a statement) differ across talkers, talker groups, and discourse contexts (Arvaniti & Garding, 2007; Tyler, 2015; Warren, 2016). For example, acoustic features of yes-no questions produced by young children and non-native talkers deviate from that of native adult talkers (Patel & Grigos, 2006).

Taken together, the existing evidence suggests that listeners encounter variable acoustic realizations of terminal contours in English. In the following sections, we report on a combination of experiments investigating the nature of the variability underlying the prosodic input and possible ways in which listeners resolve the uncertainty associated with the cue-category mapping. In Study 1, we collected production samples of utterances wherein a terminal contour is used to distinguish two intended prosodic

categories and evaluated multiple normalization methods. In Study 2, we conducted a comprehension experiment with a new group of native listeners. We used a between-participant manipulation to test whether listeners can learn to categorize *a priori* ambiguous tokens to accommodate characteristics of the cue-category mapping observed in a given context. In the general discussion, we highlight implications of these two mechanisms in phonological and processing theories of speech prosody.

STUDY 1

Native speakers of American English produced questions and statements in the form “It’s X-ing(?)”. Spectral and durational cues were extracted and examined through trained Bayesian Ideal Observer classifiers.

Methods

Participants

Fifty-nine (14 male, 45 female) undergraduate students of University of Rochester participated in exchange for monetary compensation. All participants were native speakers of American English and naive as to the purpose of the experiment.

Materials

We constructed 24 sentences in the form: *It’s X-ing*. The verbs used were all monosyllabic but varying in their syllable structures (See Appendix I).

Procedure

Each subject attended a recording session individually in a recording booth. They were instructed to read the sentence on the screen first and then produce the sentence out loud as a question or a statement according to the punctuation displayed at the end of a sentence (e.g., *It's raining.* vs. *It's raining?*). After producing each sentence, they clicked on a button on the screen to proceed to the next sentence in a self-paced manner. They produced all 24 sentences, once as a question and once as a statement, in a randomized order (48 trials in total).

Annotation

The recorded sentences were annotated by an experimenter and segmented into the following three syllables: 1) it's 2) the stressed syllable (verb stem), and 3)-ing. We extracted F0 and duration of these three segments, which are summarized in Table 1 and Fig. 1.

Table 1. Duration and F0 means and standard deviations for each segment.

	It's	X	ing
Question Duration (s)	0.237 (0.074)	0.275 (0.064)	0.298 (0.067)
Statement Duration (s)	0.241 (0.067)	0.277 (0.062)	0.303 (0.073)
Question F0 (Hz)	191 (46)	177 (46)	248 (72)
Statement F0 (Hz)	193 (48)	175 (46)	160 (51)

Figure 1. Summary of duration and pitch of produced sentences.

Hence for we treat the two contours as the question and the statement *categories*. In the literature, categories associated with the terminal contours are often defined as “high vs. low” or “rise vs. fall” (e.g., Pierrehumbert, 1980). The goal of the current investigation is to capture 1) how acoustic cues are distributed when talkers intend to distinguish a question from a statement through prosody; and 2) how listeners arrive at the two meanings through categorizing the perceived cues. Classification accuracy both in computational models and in humans is therefore evaluated with respect to the distinction between the two possible prosodic categories supporting the question and statement meanings.

Results and Discussion

We focused our analyses on the third syllable (the “-ing” part of the “It’s X-ing” sentences). Past research strongly suggests that the distinction between question and statement prosody in English is primarily, if not exclusively, signaled by cues such as f_0 and duration of an utterance final boundary tone (Patel & Grigos, 2006). This assumption is validated in our data (Figure 1 above). We analyzed f_0 and duration in four different

ways to capture their overall distributions. First, we measured raw mean f_0 and durations for the last syllable of the utterance (Figure 2-a). Second, to take into account the talker’s baseline f_0 and speech rate, we obtained F_0 and duration of the last syllable relative to those of the preceding syllable (Figure 2-b). This approach is consistent with the idea of utterance-internal normalization, where listeners interpret later-occurring acoustic information in proportion to the preceding context in the same utterance. Finally, we also used Z-score standardized f_0 and durations for the last syllable (Figure 2-c) and for the between syllable differences (2-d). These two last measures take into account each talker’s means and variances for the relevant cues, thereby placing all of the talkers on a comparable scale. In contrast to Figure 2-b, these transformations rely on having access to distributions of cue values *across* utterances produced by a given talker. We then adopted a Bayesian Ideal Observer classifier (Kleinschmidt, 2019; Kleinschmidt & Jaeger, 2015) on each measure to compute a classification accuracy. We report results of these models below (Details can be found in Appendix II).

Figure 2. Cue values for all productions. Ellipses capture 95% of the data for each category: a.): raw F0 and durations for last syllables. b.): difference scores between last syllable and previous syllable. c.) Z-scored cue values for last syllable. d.) Z-scored difference scores between last syllable and previous syllable. e.) mean difference scores for each participant's questions and statements. Each line connects a subject's question values with their statement values. f.) a subset of four talkers' productions, chosen at random.

Three clear patterns emerged. First, acoustic cues for the question and statement categories are dissociated along the dimension of f_0 (Figure 2-a, 2-b, 2-c, 2-d) but not segment duration. This result is consistent with the previous finding that a mean f_0 value, but not a mean duration value, is a primary cue to the question-statement category membership in adults' speech. Second, there is considerable ambiguity between categories. Naturally, this uncertainty resulted in relatively low accuracy scores from the classifier (66.7%) especially when applied to the raw acoustic values (Figure 2-a). The

normalizations based on either within-utterance (Figure 2-b) or -talker (Figures 2-c, and 2-d) baselines led to better classification performances (68%, 74%, and 76% respectively), which is consistent with past studies demonstrating the effectiveness of talker- (Tang et al., 2017) and gender-based (Bishop & Keating, 2012) normalization of pitch contours as well as extrinsic normalization in phoneme perception such as cCure (McMurray & Jongman, 2011). Critically, the relatively larger degree of improvement associated with the within-talker normalization (Figures 2-c and 2-d) suggests that reliable categorization of terminal contours benefit from observing multiple instances of utterances produced by the same talker.

Finally, there is a substantial degree of variability across talkers, an issue which was not fully resolved by normalization of data according to baselines of individual talkers. In Figure 2-e, we illustrate variability of F0 and duration values for individual talkers. A singular transformation function will be unlikely to permit categorical distinction between the question or statement categories. To drive this point home, in Figure 2-f, we present a random subsample of by-talker cue distributions. Although f_0 is generally used as a primary cue to differentiate the two categories, there are a great deal of individual differences in terms of its “informativeness”, or relative predictive values over category identities (see also Francis, Kaganovich, & Driscoll-Huber, 2008; Holt & Lotto, 2006; Idemaru & Holt, 2011; Kleinschmidt, 2019; Toscano & McMurray, 2010). For instance, in S20 both f_0 and duration are used as informative cues to differentiate the categories while S32 and S32 rely primarily on F0. In S54, neither of these cues seem to reliably carry information about a category identity.

In summary, the normalized data, compared to the raw acoustic data, indeed led to more accurate classification of the two underlying categories. The results corroborated the recent finding that the human comprehension system tracks prosodic information in a context- (talker-) dependent manner (Baese-Berk et al., 2014; Reinisch & Maximilian, 2015; Tang, Hamilton, & Chang, 2017). The apparent across-talker variability limits the effectiveness of a normalization method, however. The accuracy of the Bayesian Ideal Observer classifier reached only 76% after taking into account both within- and across-utterance baselines. Note that the current data are taken from a relatively homogeneous group of speakers (college students). The amount of variability would be much larger, and the categorization accuracy lower, once we broaden the scope of talkers. These results together motivate the prediction that human comprehension system will benefit from learning to accommodate the talker-dependent variations. To test this prediction, in Study 2, we conduct a comprehension experiment to investigate listeners' ability to adapt their categorization judgments according to recent exposure.

STUDY 2

We resynthesized continua spanning between prototypical instances of questions and Statements to investigate whether listeners can quickly estimate and accommodate new category distributions (in this case, accommodate a talker that produces questions or statements in a different manner) after a brief exposure to the input.

Procedure

One-hundred and eighty participants were recruited through an online platform Amazon Mechanical Turk (<https://www.mturk.com/>). Two participants were excluded for providing a uniform response (either question or statement) for all questions, leaving 178 participants that were analyzed. Participants were all self-identified native speakers of American English and received monetary compensation for their participation.

Materials

We constructed a 12-step continuum using the resynthesizing technique described in Kurumada et al. (2017). A female native speaker of American English recorded the “it’s X-ing” sentence including different verbs (cooking, booting, cooling, losing, moving, muting) with the statement and the question intonation twice each. “Cooking” was meant to be the test item and the other five items, all involving the rounded front vowel, were intended to be exposure items. These recorded stimuli were divided into three regions corresponding to three syllables (i.e., it’s |X | ing). The turning point in the f_0 contour within the “X-ing” portion was used to delineate the final two syllables. The f_0 of each region was sampled at 20 equally spaced time points, and measures from each time point were aggregated

across items to derive mean f_0 contours for statement and question contours, following Isaacs and Watson (2010). Similarly, the durations of each region were averaged across items by contour type.

Figure 3. Pitch track and durational measures of items at Steps 1–12

12-step continua were derived in the same technique used in Kurumada, Brown, and Tanenhaus (2017). Mean f_0 contours and durations were manipulated by interpolating between values within each region and manipulating f_0 and duration of each recording to match the interpolated values using the pitch-synchronous overlap-and-add algorithm implemented in Praat (Boersma & Weenink, 2008; Moulines & Charpentier, 1990). The five types of exposure items were created based on f_0 and durational values taken from those of “It’s moving”, which was done to ensure that the 12 steps of all of the exposure items had identical f_0 and durational values across the verb types. On the other hand, the 12 steps of the test item “cooking” were constructed based on recordings of “It’s cooking”. Consequently, the exact acoustic values for the exposure and the test items were distinct from each other.

Norming

The stimuli were normed by 120 participants recruited through Amazon Mechanical Turk. Twelve participants were excluded for having used the same responses for all questions, leaving 108 participants that were analyzed. Participants were all self-identified native speakers of American English and received \$1 for their participation. They were asked to listen to a talker producing either a statement or a question and make a two-alternative forced choice judgment by clicking on one of the two togglable buttons labeled as “statement” or “question”. Participants heard 12 tokens of *It’s X-ing*, each lexical item at two different continuum steps. Six lists were created by varying the nouns at each continuum step to mitigate possible spurious effects of item-specific factors. For example, one list used the verb “booting” as Steps 1 and 7, while another list used Steps 3 and 9 for the same item. Order of presentation was randomized across participants. Participants did not receive feedback about which of the two referents was the intended target.

Results of the norming study are presented in Figure 4. As expected, listeners were most likely to provide a statement interpretation for stimuli close to the prototypical falling contour (i.e., Step 1). Judgments shifted toward a question interpretation as stimuli became more similar to the prototypical rising contour (i.e., Step 12). Canonical falling tokens at Step 1 elicited the statement interpretation on 2.8% of the trials compared to 90.7% for canonical rising tokens at Step 12. The items for Step 7 were the most ambiguous (44.4% Question responses).

Figure 4. Proportions of question (as opposed to statement) responses by continuum steps in the Norming study. Error bars indicate bootstrapped 95% confidence intervals. Line was estimated using ggplot’s smooth function.

Based on this result, we decided to use the stimuli from only 11 Steps (Steps 2–12 in the Norming Study), which allowed us to have an equal number (five) of steps biased towards the Question and Statement responses respectively, with the middle step (Step 7 in the Norming Study) being ambiguous between the two categories.

Procedure

The experiment comprised three phases: a) pre-exposure test, b) exposure, c) post-exposure test. Participants were asked to judge whether a given utterance is meant to be a question or a statement and provide their answer in a two-alternative forced choice task by clicking one of the two togglable buttons labeled as “statement” or “question”. In the pre-exposure phase, participants heard 22 tokens of “it’s cooking” sampled uniformly along the 11-step continuum. They provided their answer in each trial and received no feedback.

In the exposure phase, participants were randomly assigned to one of the three conditions: Question-bias, No-bias, Statement-bias. In all the conditions, they continued

to answer 2AFC questions with feedback. All participants heard the five exposure item types six times each: three times with the statement feedback and three times with the question feedback (30 total trials). Item order was randomized. The sampling patterns of the items are illustrated in Figure 5.

Figure 5. Sampling patterns in exposure phase across three between-subject conditions.

In the *No-bias* condition, exposure items with statement feedback were sampled from Step 1, and those with question feedback from Step 11 (Figure 5b). In the *Statement-biasing* condition, exposure items with statement feedback were sampled from Steps 6, and those with question feedback from Steps 11 (Figure 5c). In contrast, in the *Question-biasing* condition, exposure items with statement feedback were sampled from Steps 1, and those with question feedback from Steps 6 (Figure 5a). Consequently, items selected from Steps 6, which had been normed to be the most ambiguous, differed in feedback across the two conditions.

The test phase was identical to the pre-exposure phase, and constant across the three between-subject conditions. Participants made 2AFC judgments for 22 tokens of *It's cooking* (with no feedback) sampled uniformly from each of the 11 steps along the

continuum. If categorization of intonation contours is modulated according to the acoustic input given in the exposure phase, exposure to the input in the Question-bias condition should result in recalibration of the categorization function such that the previously ambiguous items should be more likely to be judged as a question. On the other hand, in the Statement-bias condition, we ought to see an opposite trend. We do not expect to observe a significant change in the categorization function in the No-bias condition.

Results and Discussion

Responses in the pre- and post-exposure (test) phase are plotted in Figure 6. As predicted, in the No-bias condition, participants' judgments did not differ between pre- and post-exposure phases. This result suggests that by simply participating in this task, listeners did not shift their categorization judgments. Whatever shifts observed in the other two conditions are, therefore, attributable to the manipulations in the input given during the exposure phase. Importantly, in the Statement- and Question-biasing conditions, items from the mid region of the continuum (Steps 6-8) elicited differing interpretations. Specifically, the proportion of statement responses increased by 16.9% at Steps 6–8 in the statement-biasing condition. Similarly, the proportion of question responses increased by 18.0% at Steps 6–8 in the Question-biasing condition.

Figure 6. Proportions of question responses across conditions in Study 2. Error bars indicate bootstrapped 95% confidence intervals. Lines were fitted by using ggplot's smooth function.

We analyzed judgments with multilevel logistic regression models implemented using the `glmer` function within the `lme4` package in R (Bates, Mächler, Bolker, & Walker, 2015; R, 2010). In the statistical analysis, we compared the Statement- and Question-biasing conditions only (the No-biasing condition served as a baseline but we did not *a priori* predict any effect of exposure). We included exposure condition (Statement- vs. Question-biasing, contrast-coded with non-biasing as the comparison group), continuum steps (Steps 1–11, mean-centered, coded as a continuous variable), block (pre vs post-exposure), and their interaction as fixed effects. We used the maximum random effect structure justified by the data. When a model did not converge, we removed random slopes with the highest-order interaction terms accounting for the least variance. The final model contained only by-subject intercepts.¹

¹ Because only the bare random effects structure converged, we fitted `brms` (Bayesian Regression Models using Stan, Bürkner (2019)) models to verify that the same patterns in results hold true when tested with a more detailed random effects structure. The model included random slopes and intercepts per subject for block and continuum step. More complex random effect structures did not meaningfully improve the fit of the model. The results from this model were the same as reported above, with the exception of the main effect for block which was not supported under the new models. Details for this model are presented in Appendix III.

We found main effects of block ($\hat{\beta} = .07, p < .05$) and continuum step ($\hat{\beta} = .78, p < .001$). Critically, the interaction term between block and exposure condition was significant (for question-biasing: $\hat{\beta} = .34, p < .001$; for statement-biasing: $\hat{\beta} = -.28, p < .001$). As predicted, participants in the Question-biasing condition provided more question responses during post-exposure, and participants in the Statement-biasing condition provided fewer question responses during post-exposure.² Participants in the Question-biasing condition were more likely to provide question responses for items closer to the higher end of the continuum than participants in the Statement-biasing condition (i.e., Step 11). This observation suggests that listeners recalibrated their interpretations of *It's X-ing* to better reflect the characteristics of the recent input from exposure items.

Table 2: Summary of multilevel logistic regression model with Block (pre- vs post-exposure, sum coded with post-exposure as 1 and pre-exposure as -1), Condition (sum coded with Question-Biasing as 1, Statement-Biasing as -1, and Non-ambiguous as the comparison group), their interactions, and Continuum (centered) as fixed effects.

	Estimate ($\hat{\beta}$)	Standard Error	Z value	p value
Intercept	-0.859	0.071	-12.04	p < .001
Block	0.072	0.034	2.13	p < 0.04
Question-Biasing	0.012	0.098	0.13	p = 0.899
Statement-Biasing	-0.069	0.098	-0.70	p = 0.486
Continuum	0.775	0.017	45.36	p < .001
Block * Question-Biasing	0.336	0.048	6.96	p < .001
Block * Statement-Biasing	-0.276	0.048	-5.73	p < .001

² The categorization function of the **pre**-exposure phase in the Question-biasing condition, as compared to the No-bias and the Statement-biasing conditions, was rather shallow: only 50% of the items from Step 7 were judged to be questions. This result can be within the range of variability expected in a between-subject design such as the one adopted here. To ensure that the effect is replicable with a new group of participants, we repeated the same procedure for the Question-biasing condition, which replicated the results reported here (Appendix IV).

Consistent with our prediction, listeners learned to accommodate the distributional input in the exposure phase and provided distinct judgment patterns in the post-exposure phase. Inspection of Figure 6 suggests that the largest effect was observed in the steps where listeners find experimental stimuli *a priori* ambiguous between the two interpretations. The effect was robust, so much so that the physically identical stimuli from Step 7 elicited opposite responses in the post-exposure phase: they were judged to be questions 70% of the time in the Question-biasing condition while they were considered to be statements 28% of the time in the Statement-biasing condition). This result suggests that listeners can effectively learn to adapt their intonation interpretations in a context-dependent manner through a brief exposure (only 15 ambiguous and 15 unambiguous tokens).

General Discussion

Summary of results

The past 20 years of prosodic research has been guided by the idea that human listeners map continuous acoustic cue values to prosodic categories to infer the meaning behind them. The mapping is, however, variable and fluid across contexts. Determining how listeners arrive at the mapping *as intended by the talker* despite the surface variability constitutes one of the most outstanding questions in the literature. The current study examined two mechanisms that are expected to contribute to understanding the mechanisms: normalization and learning. Our results from Study 1 approximated “upper limits” of classification accuracy achieved through multiple context-dependent (e.g.,

utterance- or speaker-based) normalization methods. We found that talkers differ significantly in their prosodic production, and that talker-normalized values would best support accurate classification. This observation motivated the hypothesis that listeners learn mapping between acoustic cues and prosodic categories for a given talker. Study 2 demonstrated that participants were in fact able to learn distinct cue distributions over question vs. statement categories. After being exposed to only 30 tokens of utterances, participants in the Question- and Statement-biasing conditions provided opposing interpretations for the prosodically ambiguous input. We thus provided proof-of-concept evidence for the human comprehension capacity to fine-tune the cue-category mapping in prosody in a context-sensitive manner.

Notably, the categorization shifts in the two between-subject conditions in Study 2 were induced by exposure to tokens sampled from only two of the continuum steps. The overall shifts across continuum steps would not be predicted if the learning was restricted to the cue-category mapping directly experienced in the exposure phase. Rather, the results suggest that the limited exposure to discrete token samples allowed participants to learn an underlying, context- (e.g., talker-) specific, distribution of acoustic cues under the two categories. Note also that the exposure items and the test items differed in the lexical items (e.g., [exposure] “It’s {booting, moving, loading, cooling, losing, muting}” vs. [test] “It’s cooking”). The learning during the exposure phase therefore likely operated over prosodic cue-category mapping abstracted from the segmental features (e.g., phonemes) of utterances. This finding is consistent with a recent finding that a population of neurons (in the superior temporal gyrus) selectively represent speaker-normalized relative pitch, and that these neurons were distinct from those that

represent phonetic features of speech (Tang et al., 2017). The current results thus support the emergent view that prosodic information is represented and processed in a manner that is robust to the surface variability of the speech signal.

Implications for intonational speech prosody

In recent years, studies have increasingly recognized that speech prosody contains a wealth of information about socio-indexical features of talkers such as age, gender, ethnicity, language background, and social identities (Arvaniti & Garding, 2007; Breen, Kurumada, Wagner, Watson, & Yu, 2018 (and the papers therein); Clopper & Smiljanic, 2011). In a similar vein, prosody has attracted attention as a diagnostic tool for linguistic and other impairments such as the autism spectrum disorder (e.g., Jun & Bishop, 2015; McCann & Peppe, 2003; Shriberg et al., 2001). The realization that such variation, or *accents*, in prosody can be a rich source of social meaning (including talker identity), rather than a mere source of noise, has galvanized substantial excitement in the fields of sociolinguistics and experimental / clinical psychology (Cutler, 2015; Dilley, Wieland, Gamache, McAuley, & Redford, 2013; Fitzsimons et al., 2001; Green, 2002; Holliday, 2019; Podesva & Callier, 2015; Shriberg et al., 2001; Woynaroski et al., 2017 among others). At the same time, the excitement illuminated a pertinent question regarding variability in empirical approaches to intonational speech prosody (Cole, 2015; McQueen & Dilley, 2019). That is, no existing framework adequately explains how listeners can represent and process abstract, linguistically meaningful, categories (e.g., pitch accents and boundary tones) when acoustic realizations of these categories vary across talkers and contexts.

The results of the current study provide two core insights to this inquiry. First, the listeners' representational knowledge can be productively understood as distributions of acoustic cues associated with hypothesized prosodic categories. Figure 7-a represents distributions of the acoustic data collected in Study 1, characterized with distinct means and variances. According to these distributions, tokens as indicated as A and B in Figure 7-a would fall into the same category (i.e., statement) although they differ in their durational values. To compare, Figure 7-b plots hypothesized data pattern for questions and statements produced by preschoolers. Patel and Grigos (2006) demonstrated that the precise control of pitch movements matures only gradually through childhood. Hence, four-year-olds relied more on final-syllable lengthening to differentiate a question from a statement as compared to older children, resulting in miscategorization of their production by adult listeners (Patel & Brayton, 2009). The hypothesized distributional statistics of *durational* values in Figure 7-b are therefore depicted with two non-overlapping means while the pitch values are largely inseparable. Evaluated against these assumptions, Token A' and Token B' would be more reliably assigned to the two categories compared to the case of Figure 7-a. This flexible assignment is part and parcel of learning, or adaptation, that we put forward under Study 2: by adapting underlying distributional assumptions to characteristics of the input, a given token can be associated with more than one prosodic category. Put differently, listeners can maintain representations of the abstract categories while being able to map them onto acoustic realizations of speech in a context-sensitive manner.

(a) Adult data from Experiment 1 (b) Child (hypothesized) data

Figure 7. Distributions of acoustic cues and density functions associated with question vs. statement categories: (a) data from adult native speakers of English (corresponding to Panel (c) in Figure 2); (b) hypothesized data from 4-year-olds based on the observation in Patel and Grigos (2006).

Second, and relatedly, this distribution-based approach provides a new window into the mechanism underlying prosodic processing. Although differing in important details, most of the extant models of prosodic processing begin with recognizing perceived patterns of sounds (e.g., pitch movements and locations of peaks and valleys) characteristic to prosodic categories such as pitch accents (Ito & Speer, 2008; Watson, Tanenhaus, & Gunlogson, 2008; Weber, Braun, & Crocker, 2006). The recognition is immediate despite the noisy and variable acoustic realizations (Ito, Turnbull, & Speer, 2017), and the meaning is derived incrementally and, sometimes, even predictively (Bögels & Torreira, 2015; Kurumada, Brown, Bibyk, Pontillo, & Tanenhaus, 2014). It is, however, unclear what kind of pattern recognition mechanism would permit arriving at a particular category as rapidly as listeners do given that a large proportion of utterances are prosodically rather ambiguous in their category affiliations (Study 1, Figure 2).

A key to this question may be that listeners *infer* the cue-category mapping based on the distributional representations rather than finding a pitch pattern that is diagnostic of an underlying representation (for similar discussions in the domain of speech

perception, see Chodroff & Wilson, 2017; Feldman, Griffiths, & Morgan, 2009; Kleinschmidt, 2019; Kleinschmidt & Jaeger, 2015; Norris & McQueen, 2008; Norris, McQueen, & Cutler, 2015). In the case of the age-related prosodic differences between question vs. statement intonations discussed above, the expected means and variances for f_0 and duration allow listeners to derive expectations about acoustic cue values likely produced by adults (or children) under a hypothesized category. As they receive acoustic cues in the input, listeners can evaluate compare them against these “expected” configurations of cues to infer the most likely cue-category mapping given the current estimates (Farmer, Brown, & Tanenhaus, 2013). Any mismatch between the expected and observed data may then feed back into the distributional assumption, updating the means and variances for the hypothesized categories. Clearly, this is an extremely partial sketch of how the aforementioned challenges due to variability might be met. Our present purpose is chiefly to point to potential directions for exploring real-time prosodic processing system that is flexible enough to navigate the lack of invariance in the input, but sufficiently constrained to maintain abstract prosodic categories that are linked to meanings.

Let us discuss an example where the current distribution-based approach provides a clear explanation to an otherwise rather puzzling finding. Warren (2017) conducted a mouse-tracking experiment, investigating comprehension of a question vs. statement intonation contours in native listeners of New Zealand English (NZE). As in the current study, participants of Warren’s experiment heard tokens with a rising or a falling tone at the end and provided judgments as to whether they were questions or statements (e.g., “John’s mother cared for stay animals {./?}”). The critical tokens were manipulated in

terms of acoustic realizations of a diphthong that is characterized by an ongoing merger. In NZE, the diphthong in the word “care” is more likely to be produced as /ea/ by older talkers but as /ia/ by younger speakers (Gordon & MacLagan, 2001) (both variants were spoken by the same 25-year-old talker in this experiment). NZE listeners in the experiment were more likely to judge a rising terminal contour as a question when the diphthong was pronounced as /ea/ (i.e., associated with older talkers) as compared to when the diphthong was pronounced as /ia/ (i.e., associated with younger talkers). Note that this result is not readily amendable to an existing framework of prosodic processing, wherein a cue-category mapping (e.g., question vs. statement) is made primarily based on the acoustic features of the terminal contour (e.g., f_0). Warren argued that a rising intonation attracted a statement interpretation in the presence of the /ia/ variant because of listeners’ prior experiences with young talkers who are also likely users of “uptalk”, a speech style frequently involving rising declarative utterances. On the other hand, the /ea/ variant evoked the association with older talkers who are unlikely to use a rising contour for a statement category.

The current distribution-based approach concurs with this view and provides a correct prediction for the distinct judgments by NZE listeners with varying degrees of exposure to younger and older talkers. We hypothesize that NZE listeners have built underlying distributional representations for acoustic cue values associated with the question and statement categories. We also expect that the distributional information is conditioned on co-varying features such as the talker’s socio-indexical features including age and gender. In language comprehension, listeners integrate information about the talker from multiple sources (e.g., segmental phonology, voice quality, visual

information) to calibrate their expectations for acoustic cues to be used under a given prosodic category (e.g., What kinds of acoustic cues would this talker produce if the category is {question vs. statement}?). These talker-sensitive expectations influence an outcome of inference based on a perceived token of prosodic input (e.g., a rising terminal contour). An added advantage of the distributional-approach is that, with additional computational assumptions, it can make quantitative predictions about whether and how much listeners' prosodic judgments profit from keeping track of a particular socio-indexical variable of talkers (e.g., age vs. gender). An emerging body of computational work, originally developed to solve similar problems in speech perception, can be directly applied for this purpose (Kleinschmidt, 2019; Kleinschmidt, Weatherholtz, & Jaeger, 2018).

Limitations and questions for future studies

A clear limitation of the current data stems from the highly restricted number and types of utterances used in the experiments. Surely, the production data in Study 1 had only one construction ("It's X-ing") and the training and test items in Study 2 had the identical structure. The current stimuli therefore offer only a limited amount of information about to what extent the effects of talker-specific learning are generalizable across different lexical items and sentence structures. Furthermore, the current data do not offer decisive insights into the nature of the learning during the exposure. One possibility is that listeners simply learned to adjust a categorization boundary (e.g., participants in the question-biasing condition simply learned to increase the probability of the "question" response for ambiguous tokens). Another possibility is that the exposure

tokens allowed for reorganization of the internal structure of a distribution (e.g., mean and variance) exposure (Xie, Theodore, & Myers, 2016). Our current exposure items in Study 2 were sampled from one part of the continua (e.g., Steps 1 and 6 in the Question-biasing condition and Steps 6 and 11 in the Statement biasing condition), making it unclear whether and to what extent listeners actually learned underlying *distributions*.

To gain further insights, one must use a wider variety of input items, with distinct sampling patterns, to manipulate means and variances of the postulated distributions (Clayards, Tanenhaus, Aslin, & Jacobs, 2008). Furthermore, it is necessary to expand the scope of inquiry to cover a broader range of linguistic constructions that differ in their phonological and phonetic make-ups as well as their lexical and syntactic contents. Finally, to extend the current approach to include variability conditioned on discourse contexts and speaking conditions, one needs to systematically control contributions of contextual factors by introducing production paradigms that simulate language use in different situations (e.g., reading vs. speaking, casual vs. formal speech, monologue vs. dialogue). Such systematic experimentations will provide evidence for the nature of representations and processes underlying prosodic comprehension in the face of the tremendous variability in the acoustic realization of speech prosody.

Conclusion

The current study began with a puzzle: how can listeners reliably interpret meaning behind the prosodic signal when finding stable mappings between the acoustic cues and prosodic categories appears to be so difficult? We proposed that listeners maximize their chances of correctly interpreting a prosodic signal by dynamically adapting their

interpretations to their experience with, for example, a talker. The results of the two experiments demonstrated that acoustic realizations of prosodic categories exhibit a substantial amount of variability and that tracking distributional statistics of the acoustic cues associated with a context (e.g., talker identity) improves the classification accuracy.

The results of the current study critically extend an existing approach to prosodic processing, providing a testable account for how listeners can navigate the variable mapping between the acoustic cues and more abstract prosodic categories. Variability stemming from talkers and contexts can be used as a source of information about underlying distributions of acoustic cue values, which harnesses inferences over the most likely prosodic category given the observed input. These findings offer new insights into inferential processes underlying rapid and accurate prosodic processing, advancing our understanding of communication of meaning through the dynamic and variable human speech sounds.

Acknowledgments

We would like to thank Meredith Brown, Dave Kleinschmidt, and Xin Xie for providing code and assistance with R and Praat scripting; Sherwin Nourani, Manasvi Chaturvedi, and Nicole Vieyto for assistance in making materials and annotations; T. Florian Jaeger, Xin Xie, Cassandra Jacobs, and the members of the HLP Lab for providing suggestions and discussions. Research funds were provided by the provost's office at the University of Rochester.

Appendix I: Sentences produced by participants in Experiment 1, once as a statement and once as a question

1. It's closing.
2. It's ending.
3. It's falling.
4. It's freezing.
5. It's loading.
6. It's melting.
7. It's raining.
8. It's snowing.
9. It's starting.
10. It's working.
11. It's ticking.
12. It's printing.
13. It's running.
14. It's coming.
15. It's reading.
16. It's writing.
17. It's moving.
18. It's learning.
19. It's changing.
20. It's stopping.
21. It's walking.
22. It's sinking.
23. It's crashing.
24. It's cooking.

Appendix II: Classifier descriptions for Experiment 1

To assess how useful different treatments of duration and f_0 were, we trained classifiers such as the marginal models described in Kleinschmidt (2019). In Kleinschmidt (2019), models were trained to estimate the informativity and usefulness of including gender and dialect information when classifying vowels and plosives by their most relevant cues (F1xF2 for vowels, and VOT for plosives). The classifiers estimate the posterior probability distributions of cues to linguistic categories as normal distributions from a subset of the production data. It then uses these distributions to estimate the likelihood that a set of cues would belong to each of the categories. Marginal models estimated these distributions from the productions of all speakers. In other words, it treated speakers of different genders and dialects as the same. More specific models estimated additional distributions per gender or per dialect. In this case, the models have the opportunity to be more accurate if it is the case that speakers of different genders or dialects are systematically producing phonetic categories in different ways from each other.

Translating this approach to our study, we trained a classifier that would behave as the marginal models in Kleinschmidt (2019). As such, the classifiers estimated a bivariate normal distribution based on durational and f_0 cues, and from that could classify whether a sentence should be a statement or a question depending on the likelihood that a cue combination would result in category membership for each of these categories. Classifier accuracy is tested by using leave-one-out cross validation on specific lexical items: the classifier is trained on all verbs except one, and the one held-out item is used to

test the accuracy. This process is done while cycling through all the verbs, and the average of the accuracies is reported as the classifier accuracy. Thus, for each verb, a distribution is modeled by using the mean and variances of f_0 and durations of statements that do not contain that verb, and another distribution is modeled using the same parameters but taken from questions that do not contain that verb. The test data are then the cue values taken from productions that did contain the held-out verb. Each cue value pair (f_0 and duration values) is classified depending on the likelihood that that pair of values came from each distribution. Given our production data, an example could be that a syllable that was produced with an f_0 value of 300 and duration of 0.3 seconds will be classified as a question because the probability density for the question distribution is higher in that area than the probability density for the statement distribution. Similarly, a syllable produced with an f_0 value of 100 and duration of 0.3 seconds will be classified as a statement for the same reason.

To judge the importance of tracking individual talker identity, we also tested talker-specific models. These models work in the same way as the marginal models described above, except they estimate distributions for each talker. As such, this approach increases the likelihood that the classifier will be accurate in cases where there is substantial variability between talkers. One way of visualizing these models can be seen in Figure 2-e, where each talker has a question and a statement distribution fitted to their productions. These classifiers were estimated and tested with the same leave-one-out cross validation as the marginal models. Talker-specific models consistently performed better than their less-specific counterparts, achieving a 74–76% accuracy rate with the better cues.

Appendix III: BRMS Models Summary

Because only the bare random effects structure converged, we fitted BRMS (Bayesian Regression Models using Stan) models to verify that the same patterns in results hold true when tested with a more detailed random effects structure. BRMS is an R package that allows for Bayesian statistical analyses by specifying priors. However, using a uniform prior is equivalent to mixed-effects models such as those estimated via lme4. As such, results between this model and the originally intended mixed-effect model should be comparable. The model included random slopes and intercepts per subject for block and continuum step. More complex random effect structures did not meaningfully improve the fit of the model. The results from this model were the same as reported above, with the exception of the main effect for block which was not supported under the new models.

Table 3: Bayesian multilevel logistic regression model with naïve priors and maximal random effects structure justified by the data, which included random intercepts and slopes for Block and Continuum per subject.

	Estimate (β)	Estimate Error	95% CI Lower	95% CI Upper
Intercept	-1.68	0.14	-1.96	-1.41
Block	-0.03	0.09	-0.20	0.15
Question-Biasing	0.16	0.16	-0.14	0.47
Statement-Biasing	-0.26	0.16	-0.59	0.06
Continuum	1.67	0.10	1.48	1.88
Block * Question-Biasing	0.55	0.12	0.32	0.80
Block * Statement-Biasing	-0.61	0.12	-0.85	-0.37

Appendix IV: Question-Biasing Replication

Because the pattern of responses observed in the Pre-Exposure block of the Question-Biasing condition was slightly different from those observed in the other conditions, we attempted to replicate the result found in this condition. This process was done to rule-out the possibility that any change in responses between Pre- and Post-Exposure was simply because of regression to the mean because Pre-Exposure responses were more Statement-biased than expected. The replication (run with 60 subjects) led to the same pattern as the original study.

Figure 8. Summary results of replication of Question-Biasing condition. Error bars indicate bootstrapped 95% confidence intervals. Shades indicate standard error as estimated by ggplot's smooth function. Lines were fitted by using ggplot's smooth function with family = binomial.

Table 4: Summary of multilevel logistic regression model for Question-Biasing replication. Model included Block and Continuum as fixed effects, and random intercepts and slopes for Block per subject as its random effects structure.

	Estimate (β)	Standard Error	Z value	p value
Intercept	-6.466	0.306	-21.144	$p < .001$
Block	0.365	0.110	3.309	$p < .001$
Continuum	0.885	0.035	25.029	$p < .001$

References Cited

- Adank, P., Nuttall, H. E., Banks, B., & Kennedy-Higgins, D. (2015). Neural bases of accented speech perception. *Frontiers in Human Neuroscience*, 9, 1–7.
<http://doi.org/10.3389/fnhum.2015.00558>
- Adank, P., Smits, R., & van Hout, R. (2004). A comparison of vowel normalization procedures for language variation research. *The Journal of the Acoustical Society of America*, 116(5), 3099–3107. <http://doi.org/10.1121/1.1795335>
- Arnold, J. E., & Watson, D. G. (2015). Synthesizing meaning and processing approaches to prosody: Performance matters. *Language, Cognition and Neuroscience*, 30(1–2), 88–102. <http://doi.org/10.1080/01690965.2013.840733>
- Arvaniti, A. (2017). The Autosegmental-Metrical model of intonational phonology. In S. Shattuck-Hufnagel & J. Barnes (Eds.), *Prosodic theory and practice*. MIT Press.
- Arvaniti, A., & Garding, G. (2007). Dialectal variation in the rising accents of American English. *Papers in Laboratory Phonology* 9, 547–576.
- Arvaniti, A., & Garding, G. (2007). Dialectal variation in the rising accents of American English. In J. Cole & J. H. Hualde (Eds.), *Papers in Laboratory Phonology* 9 (pp. 547–576). Berlin, New York: Mouton de Gruyter.
- Baese-Berk, M. M., Heffner, C. C., Dilley, L. C., Pitt, M. A., Morrill, T. H., & McAuley, J. D. (2014). Long-term temporal tracking of speech rate affects spoken-word recognition. *Psychological Science*, 25(8), 1546–1553.
<http://doi.org/10.1177/0956797614533705>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48.

<http://doi.org/10.18637/jss.v067.i01>

- Beckman, M., & Ayers Elam, G. (1997). Guidelines for ToBI labeling. Version 3.0. Ohio State University.
- Beckman, M. E., Ayers, G. M., & Ayers Elam, G. (1994). Guidelines for ToBI labelling. Ohio State University.
- Bishop, J., & Keating, P. (2012). Perception of pitch location within a speaker's range: Fundamental frequency, voice quality and speaker sex, *The Journal of the Acoustical Society of America* 2012, 132 (2), 1100-1112.
- Boersma, P., & Weenink, D. (2008). Praat: Doing phonetics by computer (Version 5.0.26).
- Bögels, S., & Torreira, F. (2015). Listeners use intonational phrase boundaries to project turn ends in spoken interaction. *Journal of Phonetics*, 52, 46–57.
- Bolinger, D. (1989). *Intonation and its uses: Melody in grammar and discourse*. London: Arnold.
- Bolinger, D. L. (1961). Contrastive accent and contrastive stress. *Language*, 37(1), 83–96.
- Braun, B., & Ladd, D. . (2003). Prosodic correlates of contrastive and non-contrastive themes in German. *Eurospeech 2003*, 789–792.
- Breen, M., Dilley, L. C., Kraemer, J., & Gibson, E. (2012). Inter-transcriber reliability for two systems of prosodic annotation: ToBI (Tones and Break Indices) and RaP (Rhythm and Pitch). *Corpus Linguistics and Linguistic Theory*, 8(2).
- Breen, M., Kurumada, C., Wagner, M., Watson, D. G., & Yu, K. M. (2018). Introducing prosodic variability. *Laboratory Phonology: Journal of the Association for*

- Laboratory Phonology*, 9(1), 5. <http://doi.org/10.5334/labphon.142>
- Brown, M., Salverda, A. P., Dilley, L. C., & Tanenhaus, M. K. (2015). Metrical expectations from preceding prosody influence perception of lexical stress. *Journal of Experimental Psychology: Human Perception and Performance*, 41(2), 306–323.
- Brown, S. W., Newcomb, D. C., & Kahrl, K. G. (1995). Temporal-signal detection and individual differences in timing. *Perception*, 24(5), 525–538.
- Bruce, G. (1977). *Swedish word accents in sentence perspective*. Lund: Gleerups.
- Büring, D. (2003). On D-Trees, Beans, and B-Accents. *Linguistics & Philosophy*, 26(5), 511–545.
- Bürkner, P-C. (2019). Bayesian Regression Models using 'Stan' ver. 2.8.0. <https://github.com/paul-buerkner/brms>
- Calhoun, S. (2010). How does informativeness affect prosodic prominence? *Language and Cognitive Processes*, 25(7–9), 1099–1140.
- Chodroff, E., & Wilson, C. (2017). Structure in talker-specific phonetic realization: Covariation of stop consonant VOT in American English. *Journal of Phonetics*, 61, 30–47.
- Clayards, M., Tanenhaus, M. K., Aslin, R., & Jacobs, R. A. (2008). Perception of speech reflects optimal use of probabilistic speech cues. *Cognition*, 108 (3), 804–809.
- Clopper, C. G., & Smiljanic, R. (2011). Effects of gender and regional dialect on prosodic patterns in American English. *Journal of Phonetics*, 39(2), 237–245.
- Clopper, C. G., & Smiljanic, R. (2011). Effects of gender and regional dialect on prosodic patterns in American English. *Journal of Phonetics*, 39(2), 237–245.
- Cole, J. (2015). Prosody in context: A review. *Language, Cognition and Neuroscience*,

30(1–2), 1–31.

- Cole, J., & Shattuck-Hufnagel, S. (2016). New methods for prosodic transcription: capturing variability as a source of information. *Laboratory Phonology*, 7(1), 8, 1-29.
- Cooper, W. E., & Paccia-Cooper, J. (1990). *Syntax and speech*. Harvard University Press.
- Couper-Kuhlen, E. (1984). A new look at contrastive intonation. In R. J. Watts & U. Weidmann (Eds.), *Modes of interpretation: Essays presented to Ernst Leisi* (pp. 137–158). Tübingen.: Gunter Narr Verlag.
- Cutler, A. (2015). *Native listening: Language experience and the recognition of spoken words*. MIT Press.
- Cutler, A., Dahan, D., & Donselaar, W. (1997). Prosody in the comprehension of spoken language: A literature review. *Language and Speech*, 40(2), 141–201.
- Dahan, D. (2015). Prosody and language comprehension. *Wiley Interdisciplinary Reviews: Cognitive Science*, 6(5), 441–452. <http://doi.org/10.1002/wcs.1355>
- Dennison, H. Y., & Schafer, A. J. (2010). Online construction of implicature through contrastive prosody. In *Speech prosody 2010 conference*.
- Diehl, R. L., Souther, A. F., & Convis, C. L. (1980). Conditions on rate normalization in speech perception. *Attention, Perception, & Psychophysics*, 27(5), 435–443.
- Dilley, L. C. (2010). Pitch range variation in English tonal contrasts: Continuous or categorical? *Phonetica*, 67(1–2), 63–81.
- Dilley, L. C., Arjmandi, M. K., Ireland, Z., Heffner, C., & Pitt, M. A. (2016). Glottalization, reduction, and acoustic variability in function words in American English. *The Journal of the Acoustical Society of America*, 140(3114).

- Dilley, L. C., Millett, A., McAuley, J. D., & Bergeson, T. R. (2014). Phonetic variation in consonants in infant-directed and adult-directed speech: the case of regressive place assimilation in word-final alveolar stops. *Journal of Child Language*, 41(1), 155–175.
- Dilley, L. C., & Pitt, M. A. (2010). Altering context speech rate can cause words to appear or disappear. *Psychological Science*, 21(11), 1664–1670.
- Dilley, L. C., Wieland, E. a, Gamache, J. L., McAuley, J. D., & Redford, M. a. (2013). Age-related changes to spectral voice characteristics affect judgments of prosodic, segmental, and talker attributes for child and adult speech. *Journal of Speech, Language, and Hearing Research : JSLHR*, 56(1), 159–77.
- Doherty, C. P., West, W. C., Dilley, L. C., Shattuck-Hufnagel, S., & Caplan, D. (2004). Question/statement judgments: An fMRI study of intonation processing. *Human Brain Mapping*, 23(2), 85–98.
- Farmer, T., Brown, M., & Tanenhaus, M. K. (2013). Prediction, explanation, and the role of generative models in language processing [commentary]. *Behavioral and Brain Sciences*, 36, 211–212.
- Feldman, N. H., Griffiths, T. L., & Morgan, J. L. (2009). The influence of categories on perception: Explaining the perceptual magnet effect as optimal statistical inference. *Psychological Review*, 116(4), 752–782.
- Fernald, A., & Kuhl, P. (1987). Acoustic determinants of infant preference for motherese speech. *Infant Behavior and Development*, 10(3), 279–293.
- Ferreira, F. (1993). Creation of prosody during sentence production. *Psychological Review*, 100(2), 233–253.

- Fitzsimons, M., Sheahan, N., & Staunton, H. (2001). Gender and integration of acoustic dimensions of prosody: implications for clinical studies. *Brain and Language*, 78(1), 94–108.
- Flynn, N., & Foulkes, P. (2011). Comparing vowel formant normalization methods. *Proceedings of ICPHS XVII*, (August), 683–686.
- Foulkes, P., Docherty, G., & Watt, D. (2005). Phonological variation in child-directed speech. *Language*, 81(1), 177–206.
- Francis, A. L., Ciocca, V., Wong, N. K. Y., Leung, W. H. Y., & Chu, P. C. Y. (2006). Extrinsic context affects perceptual normalization of lexical tone. *Journal of the Acoustical Society of America*, 119(3), 1712–1726.
- Francis, A. L., Kaganovich, N., & Driscoll-Huber, C. (2008). Cue-specific effects of categorization training on the relative weighting of acoustic cues to consonant voicing in English. *The Journal of the Acoustical Society of America*, 124(2), 1234–1251.
- Frazier, L., Carlson, K., & Clifton, C. (2006). Prosodic phrasing is central to language comprehension. *Trends in Cognitive Sciences*, 10(6), 244–249.
- Goldsmith, J. (1976). *Autosegmental phonology*. MIT, Cambridge, MA.
- Gordon, E., & Maclagan, M. (2001). “Capturing a sound change”: A real time study over 15 years of the NEAR/SQUARE diphthong merger in New Zealand English. *Australian Journal of Linguistics*, 21(2), 215–238.
<http://doi.org/10.1080/07268600120080578>
- Grabe, E., Kochanski, G., & Coleman, J. (2005). The intonation of native accent varieties in the British Isles: Potential for miscommunication? In K. Kolaczky & J.

- Przedlacka (Eds.), *English Pronunciation Models: A Changing Scene* (pp. 311–337).
Bern, Switzerland: Peter Lang.
- Green, L. (2002). *African American English: A linguistic introduction*. Cambridge University Press.
- Gryllia, S., Baltazani, M., Arvaniti, A. (2018). The role of pragmatics and politeness in explaining prosodic variability. *Proc. 9th International Conference on Speech Prosody 2018*, 158-162.
- Hawkins, S. (2003). Roles and representations of systematic fine phonetic detail in speech understanding. *Journal of Phonetics*, 31(3), 373–405.
- Hedberg, N., Sosa, J. M., & Görgülü, E. (2017). The meaning of intonation in yes-no questions in American English: A corpus study. *Corpus Linguistics and Linguistic Theory*, 13(2), 321–368.
- Heeren, W. F. L., Bibyk, S. A., Gunlogson, C., & Tanenhaus, M. K. (2015). Asking or telling - Real-time processing of prosodically distinguished questions and statements. *Language and Speech*, 58(4), 474–501.
<http://doi.org/10.1177/0023830914564452>
- Henriksen, N. (2013). Style, prosodic variation, and the social meaning of intonation. *Journal of the International Phonetic Association*, 43(2), 153–193.
- Hirschberg, J. (2002). Communication and prosody: Functional aspects of prosody. *Speech Communication: Special Issue on Dialogue and Prosody*, 36, 31–43.
- Holliday, N. R. (2019). Variation, race, and multiracial identity in linguistic research. *Wiley Interdisciplinary Reviews: Cognitive Science*, 10(1), e1480.
- Holt, L. L., & Lotto, A. J. (2006). Cue weighting in auditory categorization: Implications for first and second language acquisition. *The Journal of the Acoustical Society of*

America, 119(5), 3059.

- Horton, W. S., Spieler, D. H., & Shriberg, E. (2010). A corpus analysis of patterns of age-related change in conversational speech. *Psychology and Aging*, 25(3), 708–13.
- Idemaru, K., & Holt, L. L. (2011). Word recognition reflects dimension-based statistical learning. *Journal of Experimental Psychology. Human Perception and Performance*, 37(6), 1939–1956.
- Isaacs, A., & Watson, D. (2010). Accent detection is a slippery slope: Direction and rate of F0 change drives listeners’ comprehension. *Language and Cognitive Processes*, 25(7), 1178–1200.
- Ito, K., & Speer, S. R. (2008). Anticipatory effects of intonation: Eye movements during instructed visual search. *Journal of Memory and Language*, 58(2), 541–573.
- Ito, K., Turnbull, R., & Speer, S. R. (2017). Allophonic tunes of contrast: Lab and spontaneous speech lead to equivalent fixation responses in museum visitors. *Laboratory Phonology*, 8(1).
- Johnson, K. (2005). Speaker normalization in speech perception. D. B. Pisoni and R. E. Remez (Eds). *The Handbook of Speech Perception*.
- Jun, S.-A., & Bishop, J. (2015). Priming implicit prosody: Prosodic boundaries and individual differences. *Language and Speech*, 58(4), 459–473.
- Kleinschmidt, D. F. (2019). Structure in talker variability: How much is there and how much can it help? *Language, Cognition and Neuroscience*, 34(1), 43–68.
<http://doi.org/10.1080/23273798.2018.1500698>
- Kleinschmidt, D. F., & Jaeger, T. F. (2015). Robust speech perception: Recognize the familiar, generalize to the similar, and adapt to the novel. *Psychological Review*,

122(2), 148–203.

Kleinschmidt, D. F., Weatherholtz, K., & Florian Jaeger, T. (2018). Sociolinguistic perception as inference under uncertainty. *Topics in Cognitive Science*, 10(4), 818–834.

Krahmer, E., & Swerts, M. (2001). On the alleged existence of contrastive accents. *Speech Communication*, 34(4), 391–405.

Kurumada, C., Brown, M., & Tanenhaus, M. K. (2014). Is it or isn't it?: Listeners make rapid use of prosody to infer speaker meanings. *Cognition*, 133 (2), 335–342.

Kurumada, C., Brown, M., & Tanenhaus, M. K. (2017). Effects of distributional information on categorization of prosodic contours. *Psychological Bulletin and Review*.

Ladd, D. (1996). *Intonational phonology*. (Cambridge Studies in Linguistics 79). Cambridge University Press.

Ladd, D. R. (2008). *Intonational phonology* (2nd ed.). Cambridge University Press.

Ladd, D. R., & Morton, R. (1997). The perception of intonational emphasis: Continuous or categorical? *Journal of Phonetics*, 25(3), 313–342.

Leather, J. (1983). Speaker normalization in perception of lexical tone. *Journal of Phonetics*, 11, 373–382.

Lee, C.-Y. (2009). Identifying isolated, multispeaker Mandarin tones from brief acoustic input: A perceptual and acoustic study. *Journal of the Acoustical Society of America*, 125, 1125–1137.

Lennes, M., Stevanovic, M., Aalto, D., & Palo, P. (2016). Comparing pitch distributions using Praat and R. *Phonetician*, 111, 35–53.

- Lobanov, B. M. (1971). Classification of Russian vowels spoken by different speakers. *The Journal of the Acoustical Society of America*, 49(2B), 606–608.
- McCann, J., & Peppe, S. (2003). Prosody in autism spectrum disorders: A critical review. *International Journal of Language & Communication Disorders*, 38(4), 325–350.
- McMurray, B., & Jongman, A. (2012). What information is necessary for speech categorization? Harnessing variability in the speech signal by integrating cues computed relative to expectations, *Psychological Review*, 118(2), 219–246.
- McQueen, J. M., & Dilley, L. C. (n.d.). Chapter 34. Prosody and spoken-word recognition. In C. Gussenhoven & A. Chen (Eds.), *The Oxford Handbook of Language Prosody*.
- Monahan, P. J., & Idsardi, W. J. (2010). Auditory sensitivity to formant ratios: Toward an account of vowel normalization. *Language and Cognitive Processes*, 25(6), 808–839.
- Moore, C. B., & Jongman, A. (1997). Speaker normalization in the perception of Mandarin Chinese tones. *Journal of the Acoustical Society of America*, 102, 1864–1877.
- Morrill, T. H., Dilley, L. C., McAuley, J. D., & Pitt, M. A. (2014). Distal rhythm influences whether or not listeners hear a word in continuous speech: Support for a perceptual grouping hypothesis. *Cognition*, 131(1), 69–74.
- Moulines, E., & Charpentier, F. (1990). Pitch-synchronous waveform processing techniques for text-to-speech synthesis using diphones. *Speech Communication*, 9(5–6), 453–467.
- Newman, R. S., & Sawusch, J. R. (1996). Perceptual normalization for speaking rate:

- effects of temporal distance. *Perception & Psychophysics*, 58(4), 540–560.
- Newman, R. S., & Sawusch, J. R. (2009). Perceptual normalization for speaking rate III: Effects of the rate of one voice on perception of another. *Journal of Phonetics*, 37, 46–65.
- Norris, D., & McQueen, J. M. (2008). Shortlist B: A Bayesian model of continuous speech recognition. *Psychological Review*, 115(2), 357–95.
- Norris, D., McQueen, J. M., & Cutler, A. (2003). Perceptual learning in speech. *Cognitive Psychology*, 47(2), 204–238.
- Norris, D., McQueen, J. M., & Cutler, A. (2015). Prediction, Bayesian inference and feedback in speech recognition. *Language, Cognition and Neuroscience*, 3798(December 2015), 1–15.
- Patel, R., & Brayton, J. T. (2009). Identifying prosodic contrasts in utterances produced by 4, 7, and 11 year old children. *Journal of Speech, Language, and Hearing Research*, (June), 790–801.
- Patel, R., & Grigos, M. I. (2006). Acoustic characterization of the question-statement contrast in 4, 7 and 11 year-old children. *Speech Communication*, 48(10), 1308–1318. <http://doi.org/10.1016/j.specom.2006.06.007>
- Pierrehumbert, J. (1980). *The phonology and phonetics of English intonation*. MIT, Cambridge, MA.
- Pierrehumbert, J. (2001). Exemplar dynamics: Word frequency, lenition and contrast. In J. Bybee & P. Hopper (Eds.), *In Frequency and the Emergence of Linguistic Structure* (pp. 137–157). John Benjamins.
- Pierrehumbert, J., & Hirschberg, J. (1990). The meaning of intonational contours in the

- interpretation of discourse. In P. R. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in communication* (pp. 271–311).
- Pisoni, D. B. (1977). Some thoughts on “normalization” in speech perception. In D. Johnson & J. W. Mullennix (Eds.), *Talker variability in speech processing*. San Diego, CA: Academic Press.
- Pitrelli, J. F., Beckman, M., & Hirschberg, J. (1994). Evaluation of prosodic transcription labeling reliability in the ToBI framework. In *Proceedings of the International Conference on Spoken Language Processing* (pp. 123–126).
- Podesva, R., & Callier, P. (2015). Voice quality and identity. *Annual Review of Applied Linguistics*, 35, 173–194.
- Price, P. J., Ostendorf, M., Shattuck-Hufnagel, S., & Fong, C. (1991). The use of prosody in syntactic disambiguation. *Journal of Acoustic Society of America*, 90(6), 2956–2970.
- R. (2010). Development Core Team, R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria.
- Reinisch, E. V. A., & Maximilian, L. (2015). Speaker-specific processing and local context information : The case of speaking rate. *Applied Psycholinguistics*, 37, 1–19.
- Roettger, T. B., & Franke, M. (under review). Evidential strength of intonational cues and rational adaptation to (un-)reliable intonation.
- Saindon, M. R., Trehub, S. E., Schellenberg, E. G., & Van Lieshout, P. (2017). When is a question a question for children and adults? *Language Learning and Development*, 13(3), 274–285.
- Samuel, A. G., & Kraljic, T. (2009). Perceptual learning for speech. *Attention*,

- Perception, & Psychophysics*, 71(6), 1207–1218.
- Sawusch, J. R., & Newman, R. S. (2000). Perceptual normalization for speaking rate II: Effects of signal discontinuities. *Perception and Psychophysics*, 62(2), 285–300.
- Schafer, A., Speer, S. R., Warren, P., & White, D. (2000). Intonational Disambiguation in Sentence Production and Comprehension. *Journal of Psycholinguistic Research*, 29(2), 169–182.
- Schweitzer, K. (2012). *Frequency effects on pitch accents: Towards an exemplar-theoretic approach to intonation*. Universität Stuttgart.
- Shattuck-Hufnagel, S., Dilley, L., Veilleux, N., Brugos, A., & Speer, R. (2004). F0 peaks and valleys aligned with non-prominent syllables can influence perceived prominence in adjacent syllables. *Proceedings of Speech Prosody*, 705–708.
- Shriberg, L., Paul, R., McSweeney, J., Klin, A., Cohen, D., & Volkmar, F. (2001). Speech and prosody characteristics of adolescents and adults with high-functioning autism and Asperger syndrome. *Journal of Speech, Language, & Hearing Research*, 44, 1097–1115.
- Studdert-Kennedy, M., & Hadding., K. (1973). Auditory and linguistic processes in the perception of intonation contours. *Language and Speech*, 16(4), 293–313.
- Swerts, M., & Geluykens, R. (1994). Prosody as a marker of information flow in spoken discourse. *Language and Speech*, 37, 21–43.
- Syrdal, A. K., & McGory, J. (2000). Inter-transcriber reliability of ToBI prosodic labeling. *Proceedings of the ICSLP-20003*, 3, 235–238.
- Tang, C., Hamilton, L. S., & Chang, E. F. (2017). Intonational speech prosody encoding in the human auditory cortex. *Science*, 801(August), 797–801.

- Titze, I. R. (1989). Physiologic and acoustic differences between male and female voices. *The Journal of the Acoustical Society of America*, 85(4), 1699–1707.
- Tomlinson, J. M., Gotzner, N., & Bott, L. (2017). Intonation and pragmatic enrichment: How intonation constrains ad hoc scalar inferences. *Language and Speech*, 60(2), 200–223.
- Toscano, J. C., & McMurray, B. (2010). Cue integration with categories: Weighting acoustic cues in speech using unsupervised learning and distributional statistics. *Cognitive Science*, 34(3), 434–464.
- Tyler, J. C. (2015). Expanding and mapping the indexical field: Rising pitch, the uptalk stereotype, and perceptual variation. *Journal of English Linguistics*, 43(4), 284–310.
- Venditti, J. J., & Hirschberg, J. (2003). Intonation and discourse processing. *International Congress of Phonetic Sciences (ICPhS)*, 1–9. Retrieved from
- Wagner, M., & Watson, D. G. (2010). Experimental and theoretical advances in prosody: A review. *Language and Cognitive Processes*, 25(7–9), 905–945.
- Ward, G., & Hirschberg, J. (1985). Implicating uncertainty: The pragmatics of fall-rise intonation. *Language*, 61(4), 747–776.
- Warren, P. (2016). *Uptalk: The phenomenon of rising intonation*. Cambridge University Press.
- Warren, P. (2017). The interpretation of prosodic variability in the context of accompanying sociophonetic cues. *Laboratory Phonology*, 8(1), 1–21.
- Watson, D. G., Tanenhaus, M. K., & Gunlogson, C. A. (2008). Interpreting pitch accents in on-line comprehension: H* vs. L+H*. *Cognitive Science*, 32, 1232–1244.
- Weatherholtz, K., & Jaeger, T. F. (2016). Speech perception and generalization across

- talkers and accents. In *Oxford Research Encyclopedia of Linguistics*. (p. <http://oxfordre.com/linguistics/view/10.1093/acref>). Oxford University Press.
- <http://doi.org/10.1093/acrefore/9780199384655.013.95>
- Weber, A., Braun, B., & Crocker, M. W. (2006). Finding referents in time: Eye-tracking evidence for the role of contrastive accents. *Language and Speech*, 49(3), 367–392.
- Woynaroski, T., Oller, D. K., Keceli-Kaysili, B., Xu, D., Richards, J. A., Gilkerson, J., ... Yoder, P. (2017). The stability and validity of automated vocal analysis in preverbal preschoolers with autism spectrum disorder. *Autism Res*, 10(3), 508–519.
- Xie, X., Theodore, R. M., & Myers, E. B. (2016). More than a boundary shift: Perceptual adaptation to foreign-accented speech reshapes the internal structure of phonetic categories. *Journal of Experimental Psychology: Human Perception and Performance*, 43(1), 206–217.
- Yoon, T.-J., Chavarria, S., Cole, J., & Hasegawa-Johnson, M. (2004). Intertranscriber reliability of prosodic labeling on telephone conversation using ToBI. In *Interspeech 2004 Conference* (pp. 2729–2732). Jeju Island, Korea.
- Zipp, L., & Dellwo, V. (2011). Read speech normalization (RSN): A method to study prosodic variability in spontaneous speech. *Proc. XVIIth ICPhS Hong Kong*, (August), 2328–2331.