Running head: INDEXICAL CUES AND DISTRIBUTIONAL LEARNING

The role of indexical cues in the distributional learning of phonetic categories

Masaki Noguchi and Carla L. Hudson Kam

The University of British Columbia

Keywords: phonetic category learning, distributional learning, talker variability

Abstract

Theories of the distributional learning of phonetic categories assume that input provides reliable distributional cues for the categorization of speech sounds. In the real world, however, not all talkers produce exactly the same distributions of speech sounds, and the talker-dependent variation may undermine the reliability of the distributional cues. In this study, we investigated how learners might overcome talker-dependent variation. Specifically, we tested whether adults can learn two phonetic categories from input in which talker-dependent variation introduces potential ambiguities into the categorization of speech sounds. The results suggest that they can overcome this kind of ambiguity by using indexical information (i.e., the identity of talkers).

Introduction

Numerous studies have demonstrated that humans can learn aspects of sound structure such as phonetic features (e.g., voicing) and segments (e.g., [d] and [t]) from the frequency distribution of speech sounds in acoustic space, referred to as distributional learning of phonetic categories (Maye, 2000; Maye, Werker, & Gerken, 2002). For example, Maye et al. (2002) exposed two groups of English-learning infants to sounds taken from a continuum between the prevoiced [d] and voiceless unaspirated [t]. For one group, the frequency distribution of the sounds was bimodal. For the other group, the frequency distribution was unimodal. After exposure, infants who were exposed to the bimodal distribution discriminated [d] and [t], suggesting that they learned to classify these sounds into two categories. In contrast, infants who were exposed to the unimodal distribution failed to discriminate [d] and [t], suggesting that they learned to classify these sounds into a single category.

A basic assumption of the distributional learning hypothesis is that input provides reliable distributional cues for the categorization of speech sounds. In experiments, when the distributions are constructed and well controlled, the assumption holds. However, it may not be valid when it comes to natural language input (see, e.g., Bion, Miyazawa, Kikuchi, & Mazuka, 2013). In the real world, naturally produced speech shows a great deal of variation that can affect the reliability of distributional information. For example, vowel sounds show different distributions in acoustic space depending on the dialect of the talker (Hagiwara, 1997). However there is also variation within dialects, correlating with things like gender and age of the talker (Hillenbrand, Getty, Clark, & Wheeler, 1995). Complicating things further, this sort of variation may even occur at the level of individual talkers: when different talkers produce the same utterance, the acoustic realizations of their productions are never the same. Importantly, these

variations may undermine the reliability of distributional cues for the learning of phonetic categories. For instance, when learners are exposed to input produced by two talkers who speak different dialects and two different phonetic categories produced by these talkers show a significant overlap in acoustic space, the current distributional learning hypothesis would predict that learners may end up classifying the tokens of these categories into a single category.

In this study, we hypothesize that learners can overcome the ambiguity in input that comes from talker-dependent variation by tracking indexical information in the acoustic signal and linking it to distributional analyses. The speech signal carries different kinds of indexical information: age, gender, sexual orientation, birthplace, socio-economic status, and individual talker identity. Studies have demonstrated that adult listeners are sensitive to these talker-specific idiosyncrasies in the acoustic signal and use their knowledge about talker identity to facilitate speech processing. For example, being familiar with a talker's voice significantly facilitates various speech processing tasks such as word recognition (Mullennix, Pisoni, & Martin, 1989), word recall (Palmeri, Goldinger, & Pisoni, 1993), phonetic categorization (; Mullennix & Pisoni, 1990), and speech shadowing (Goldinger, 1996). Based on these observations, the exemplar-based model of speech perception assumes that indexical information is an integral part of linguistic knowledge (Goldinger, 1996; Johnson, 1997). In this model, linguistic knowledge is built upon the accumulation of episodic memory traces or exemplars, and each exemplar is tied to linguistic as well as indexical information. Therefore, it is possible that when learners build their knowledge about phonetic categories they integrate information about the identity of talkers, and so are able to learn to classify speech sounds according to the distributions of individual talkers. This would enable learners to overcome the ambiguity in the input introduced

by multiple talkers by processing distributional information relative to each talker's distribution rather than relying on the overall and misleading distributional across talkers.

In this study, we tested whether adults can learn two phonetic categories from input in which talker-dependent variation in the acoustic signal introduces ambiguity into the categorization of speech sounds. We exposed participants to input produced by multiple talkers. Each talker provided reliable distributional cues for the learning of two categories – specifically, a bimodal distribution of speech sounds – but the locations of the frequency peaks varied from talker to talker such that the overall distribution across talkers was unimodal. If participants only learn the overall distribution, they should end up with a single category. By contrast, if participants are able to use indexical information such that they process distributional information in a talker-specific manner, they should be able to learn that each talker produces a bimodal distribution and, therefore, end up with two categories. We also examined whether participants formed separate category structures for each talker, that is, whether they remembered the specifics of each speaker's phone distributions.

To examine these questions, we tested the learning of consonantal categories based on a duration contrast, namely, short versus long fricatives ([f] vs. [f:] and [s] vs. [s:]). In languages like Japanese (e.g., Labrune, 2012), segmental duration is contrastive-- differences in segmental duration can change the meaning of words -- whereas in English, the native language of our participants, it is not. Importantly, while studies have demonstrated that native English speakers have difficulty in perceiving differences in segmental duration, they also show that training can improve perception (Hirata, 2004; Hirata, Whitehurst, & Cullings, 2007). Thus, it is a contrast

that would have to be learned by our participants. Moreover, segmental duration categories have previously been shown to be learnable from distributional information (Pajak, 2012).¹

Another reason for using the duration contrast is that the perception of segmental duration is less affected by talker variation than other kinds of contrasts. This is crucial for testing whether learners are able to track talker-specific distribution of speech sounds in input. The distributional learning hypothesis assumes that learners are able to use the distribution of speech sounds, or the distribution of their acoustic values, as a cue for the categorization of speech sounds, but if talker variation significantly affects the perception of the acoustic values at issue, talker variation becomes a confounding factor rather than an experimental factor. It is well established that the perception of speech sounds is significantly affected by context. For example, the categorization of consonants that differ in the place of articulation, which is cued by formant transition, is affected by the formant values of the adjacent sounds (e.g., Lindblom and Studdert-Kennedy, 1967; Mann, 1980; Mann and Repp, 1980; Repp and Mann, 1981; Repp, 1981; Summerfield, 1975). Since talker variation is strongly correlated with differences in spectral values, the perception of phonetic contrasts that are based on differences in spectral values (e.g., place contrasts in consonants) is more strongly affected by talker variation than phonetic contrasts that are based on differences in non-spectral values (e.g., duration contrasts).

This does not mean that the perception of segmental duration is not affected by talker variation -

¹ Although Pajak (2012) claimed that her participants learned something about fricative length from distributional information, the methodology used makes her results somewhat unclear. Participants in her bimodal condition performed at chance level, and participants in the unimodal condition showed a significant bias towards *same* responses. Pajak interpreted the results as indicating that participants in the unimodal condition learned to classify test stimuli into a single category and that participants in the bimodal condition failed to learn anything. However, since her experiments did not assess baseline performance of naive English speakers, it is not entirely clear which group learned what, if anything.

it can be (see Cumming, 2011)² - just that it is less affected than other types of contrasts and so is a better choice than other contrasts for this investigation.

Our experiment was a modified version of the design used by Maye et al. (2002). We exposed adult English speakers to speech sounds drawn from continua between two short and long fricatives ([f] vs. [f:] and [s] vs. [s:]) produced by three talkers. The frequency of the sounds produced by each talker was manipulated so that the frequency distribution was bimodal. This means that each talker provided reliable distributional cues for the learning of a phonetic contrast between short and long fricatives. In one condition (Overlap condition), the location of the distributional peaks was the same across talkers, and thus the overall distribution was bimodal. In this condition, both the talker-specific and overall distributions provided reliable cues for the learning of the duration contrast. In the other condition (Offset condition), the location of the peaks varied across talkers such that the overall distribution was unimodal. In this condition the talker-specific distributions provided reliable cues for the learning of the duration contrast, but the overall distribution did not. We also included a control condition in which participants did not get any exposure to the sounds prior to being tested to assess native English speakers' baseline sensitivity to acoustic differences between short and long fricatives. Participants in all three conditions did the same AX discrimination test post-exposure.

We expected that participants in the overlap condition would be able to learn two duration categories based on the bimodal distribution of the fricative sounds, and thus that they would show better sensitivity to acoustic differences between short and long fricatives than

² It has been demonstrated that pitch difference affects the perception of segmental duration: dynamic F0 contours are perceived to be longer than simple F0 contours (Cumming, 2011). Even though the fricative sounds used in this experiment are voiceless and have no F0 component, pitch differences in adjacent vowels may affect the perception of the overall duration of the stimuli and thus indirectly affect the perception of the duration of the voiceless fricatives.

participants in the control condition. If participants in the offset condition are able to track talker-specific distributional information in the input, they should also learn two duration categories and so show enhanced sensitivity to differences between short and long fricatives. If, however, participants in the offset condition fail to track talker-specific distributional information, they should learn a single duration category based on the overall unimodal distribution. We also expect that if participants form separate categories for each talker, participants in the offset condition should show a talker-specificity effect; they should perform better in the discrimination task when test stimuli produced by a talker are the canonical tokens of the short and long categories produced by the same talker compared to when test stimuli produced by a talker are not their canonical tokens.

Methods

Participants

Eighty-six native English speakers participated in the experiment. All self-reported that English was their first and dominant language. Many of the participants were multilingual, but none was familiar with any language in which segmental duration is contrastive. Participants received \$10 or course credit for their participation.

Figure 1a (top) and 1b (bottom). Frequency distributions of fricative sounds in the overlap (top) and offset conditions (bottom).

Experimental manipulation

Participants were randomly assigned to one of three conditions: overlap (N=29), offset (N=28), and control (N=29). In the overlap condition participants heard three talkers, all of whom produced exactly the same variation, as shown in Figure 1a. Specifically, the three talkers, labeled Mid 1, Mid 2, and Mid 3, all produced fricatives that ranged in duration from 140-280 ms. For all three talkers fricatives with durations of 160 ms and 260 ms occurred with a much higher frequency than fricatives at other points along the duration continuum. These two points (160 ms and 260 ms) were the bimodal peaks in the distribution. Notice that these points are peaks both for the three individual talkers and for the overall summed distribution, that is, the bimodal distributional peaks overlap exactly for the three talkers.

In the offset condition, shown in Figure 1b, the distributions produced by each of the three talkers differed – they were offset from each other. Each distribution had the same shape as the other two talkers' distributions in this condition; however, the location of the peaks on the duration continuum differed for each talker. The Short talker's durations ranged from 100-240 ms, the Mid talker's from 140-280 ms, and the Long talker's from 180-320 ms. Note that the mid talker's frequency peaks are the same as those of all three talkers in the overlap condition. Importantly, all three talkers' distributions had the same shape as the three talkers in the overlap condition, and each talker's two frequency peaks were separated by 100 ms. That is, each talker evinces a bimodal distribution in their speech. However, in the offset condition the locations of the peaks do not overlap, they are offset from each other. Importantly, the specifics of the three different frequency distributions produce a unimodal overall distribution in the offset condition (see Figure 1b). The third condition was the control condition. Participants in the control condition did not receive any exposure.

Stimuli

The critical stimuli were exemplars of the bisyllabic strings /ifa/ and /asa/, where /f/ and /s/ just mean labiodental and alveolar fricative respectively. Participants were exposed to variants of the critical syllables drawn from the duration continuum, and exemplars from different points along the duration continuum occurred with varying frequencies in the input stimuli, with the specifics depending on condition as described above. Participants in the two experimental conditions heard stimuli produced by three distinct talkers, two males and one female. In the overlap condition the three talkers produced the same distributions, while in the offset condition each talker produced a frequency distribution anchored at a distinct set of locations on the continuum. Note that the range of the duration and the nature of the frequency distribution produced by every talker in both conditions was exactly the same.

There were also filler stimulus items: /ira/, /iʔa/, /atʃa/, and /adʒa/. For these items, the duration of the intervocalic consonant always matched the most frequent duration associated with the short fricative for that specific talker. In the offset condition this meant that the intervocalic consonants differed by speaker, but in a very consistent way. In the overlap condition there was no variation, as all three speakers had the same short fricative duration.

Stimuli creation

The stimuli used in the study were created from exemplars of /ifa/ and /asa/ produced by three native speakers of English. The duration of the segments was manipulated using the PSOLA method in Praat (Boersma & Weenink, 2001; Moulines & Charpentier, 1990). The duration of the first and second vowels was changed to 150 and 200 ms respectively, and the duration of the intervocalic fricatives was changed to make two continua for each speaker, one for each fricative. The filler stimuli were manipulated in similar ways: the duration of the first and second

vowels was changed to 150 and 200 ms, and the duration of the intervocalic consonants was changed to match the canonical exemplar of the short category for that talker, 120 ms for the short talker, 160 ms for the mid talker, and 200 ms for the long talker. In the overlap condition the intervocalic consonant duration for filler stimuli was 160 ms, which is the duration that occurs most frequently with the short fricatives for all three talkers.

In the offset condition, in order to ensure that any effects were not due to specifics of the voice, the mapping from specific talker to talker duration (short, mid, long) was varied in three sub-conditions (see Table 1).

Table 1

Sub-condition	Talker-duration		
	Short	Mid	Long
Offset Sub1	Male 1	Female	Male 2
Offset Sub2	Male 2	Male 1	Female
Offset Sub3	Female	Male 2	Male 1

Stimuli were arranged into blocks of 192 items presented in a random order, for a total time of 15 minutes. There were 96 critical stimuli and 96 filler stimuli.

AX discrimination test

Participants' sensitivity to acoustic differences between short and long fricatives was tested using an AX discrimination task. In Different trials, participants compared stimuli with short fricatives to stimuli with long fricatives ([f]/[s] at 160 ms vs. [f:]/[s:] at 260 ms) produced by the same talker. These same durations were tested for each talker. For participants in the overlap condition, the test stimuli were the canonical exemplars of short and long categories, that is, they were the critical stimuli with the highest frequencies in input. For participants in the offset

condition, these were the canonical exemplars of short and long categories produced by the mid talker (called Mid-Match trials), but for the short talker they were longer than the canonical exemplars (Short-mismatch trials), and for the long talker they were shorter than the canonical exemplars (Long-mismatch trials). In Same trials, participants compared two identical tokens of stimuli with either short or long fricatives. In order to make the focus of the study less obvious, filler trials were included, in which participants compared two different filler stimuli (/ira/ vs. /iʔa/, /atʃa/ vs. /adʒa/), or two identical tokens of a single filler stimulus.

There were 48 test trials (24 different and 24 same trials) and 48 filler trials (24 different and 24 same trials). As described above, in the offset condition, there were three subtypes of different trials, short-mismatch, match, and long-mismatch; there were eight trials for each subtype. Trials had an ISI of 750 ms.

Procedure

Participants were run individually in the Language and Learning Lab at the University of British Columbia. After reading and signing the consent form, for participants in the two experimental conditions, exposure began. During exposure, participants heard the stimuli block twice. After hearing each stimulus item, participants pressed the space bar to proceed to the following item. The key pressing action was included to keep participants attending to exposure stimuli. The exposure phase lasted about 15 minutes. After completing the exposure phase, participants did an AX discrimination test. In the test phase, participants had a maximum of five seconds to respond in each trial, and the trial was terminated when they recorded a response. New trials began 2 seconds after a response. Participants in the control condition skipped the exposure phase and went straight to testing. After completing the test, participants filled in a language background questionnaire.

Results

Responses to test trials in the AX discrimination test were converted into d' scores (Macmillan & Creelman, 2005). Figure 2 shows the mean d' scores by condition (control, overlap, offset) and fricative contrast ([s] vs. [s:], [f] vs. [f:]); error bars represent SE. A repeated-measures ANOVA was conducted on d' scores with condition as a between-participant variable and fricative contrast as a within-participant variable. There were significant main effects of condition [$F(2, 83) = 8.199, p < 0.001$] and fricative contrast [$F(1, 83) = 13.776, p < 0.001$], but there was no significant interaction between condition and fricative contrast [$F(2, 83) = 0.276, p < 0.759$].

Figure 2. Mean d' scores by fricative contrast and condition (Error bar represents standard error).

The main effect of fricative contrast reflects the fact that participants across conditions obtained higher mean d' scores with the alveolar fricatives ($M = 1.35$, $SD = 1.13$) than with the labiodental fricatives ($M = 1.14$, $SD = 1.09$). With condition, post-hoc pairwise comparisons indicated that the mean d' score of the control condition ($M = 0.65$, $SD = 0.72$) was significantly lower than that of the overlap condition ($M = 1.62$, $SD = 1.11$); $t(83) = 3.72$, $p = 0.001$, and that of the offset condition ($M = 1.48$, $SD = 1.18$); $t(83) = 3.23$, $p = 0.003$, but there was no significant difference between the overlap and offset conditions; $t(83) = 0.51$, $p = 0.608$. Thus, both experimental conditions showed improved sensitivity to the duration contrast as compared to the baseline sensitivity seen in the control participants, and this sensitivity did not differ according to condition; participants who received the input with an overall unimodal distribution still learned the new categories, suggesting they were tracking frequencies by talker.

We conducted further analyses on the data from participants in the overlap and offset conditions separately to examine two different questions. First we tested the effect of individual talker identity. Studies have shown wide-spread effects of individual voice characteristics on speech perception and learning (see Logan, Lively, & Pisoni, 1991; Bradlow, Torretta, & Pisoni, 1996). Although voice characteristics are not the focus of this study, they are a factor that might influence the talker-specific learning we are interested in, so it is important to ensure that any talker effects in our condition of interest (the offset condition) are not simply due to voice differences. We assessed this by examining potential differences in discrimination performance by talker for participants in the overlap condition, as they were exposed to three talkers who all produced the same bimodal distribution as each other, and thus, any differences in performance by talker in these participants could not be due to differences in other aspects of the input associated with talker.

Figure 3. Mean d' scores by fricative contrast and talker identity in the overlap condition (Error bar represents standard error).

Figure 3 shows the mean d' scores by fricative contrast and talker identity for participants in the overlap condition. A repeated-measures ANOVA was conducted with fricative contrast and talker identity (female, male 1, male 2) as within-participant variables. There were no significant main effects of fricative contrast [$F(1, 27) = 3.339, p = 0.079$] or talker identity [$F(2, 54) = 0.366, p = 0.695$], and there was no significant interaction between fricative and talker identity [$F(2, 54) = 0.332, p = 0.719$].

Figure 4. Mean d' scores by fricative contrast and talker match in the offset condition (Error bar represents standard error).

Second, we asked whether participants had learned specific distributions associated with each talker in the offset condition. Recall that the test stimuli that were compared on each trial consisted of exemplars at 160 and 260 ms. In the offset condition, these points along the continuum were the canonical exemplars of short and long categories for the mid talker (mid-match), but they were longer than the canonical exemplars of short and long categories for the short talker (short-mismatch) and shorter than the canonical exemplars of short and long categories for the long talker (long-mismatch). Figure 4 shows the mean d' scores by fricatives and talker match. A repeated-measures ANOVA was conducted with fricative contrast and talker

match (short-mismatch, mid-match, long-mismatch) as within-participant variables. There were no significant main effects of fricative [$F(1, 28) = 3.048, p = 0.0918$] or talker match [$F(2, 56) = 1.284, p = 0.285$], and no significant interaction between fricative contrast and talker match [$F(2, 56) = 1.505, p = 0.231$].

Discussion

We tested whether adults can learn two phonetic categories from input in which talker-dependent variation in the acoustic signal introduces potential ambiguity into the categorization of speech sounds. Participants in two experimental conditions were exposed to input produced by multiple talkers. Each talker provided reliable distributional cues for the learning of two categories (i.e., a bimodal distribution of speech sounds), but the locations of the frequency peaks varied from talker to talker such that overall distribution across talkers was unimodal in one condition and bimodal in another. Their post-exposure sensitivity was compared to that of another group of participants who were tested without any exposure.

The data show that participants in the two experimental conditions were more sensitive to the novel contrast than participants in the control condition, and that there were no differences between participants in the two experimental conditions. It seems then that participants in the overlap and offset conditions learned the novel phonetic categories based on durational differences. In the overlap condition, the frequency distribution produced by each of three talkers was bimodal, and the overall frequency distribution across three talkers was also bimodal. Therefore, it is not surprising that participants in the overlap condition learned to distinguish the short and long fricatives. In the input used in the offset condition, by contrast, while the frequency distribution produced by each of three talkers was bimodal, the overall frequency distribution was unimodal. This talker-dependent variation could have introduced ambiguity into

the categorization of the fricative sounds. Despite this, participants in the offset condition seem to have succeeded in learning the duration contrast. This suggests that they were able to track distributions by talker. In other words, they were able to use indexical information in their distributional computations.

The finding that talker match did not have an effect on sensitivity for participants in the offset condition suggests that they were not, however, learning three distinct distributions. Rather, they acquired some generalized knowledge about the novel phonetic contrast. This seems surprising in light of the large amount of evidence showing that familiarity with talkers' voice significantly affects speech perception (talker-specificity effects: Craik & Kirsner, 1974; Creelman, 1957; Goldinger, 1996; Goldinger et al., 1991; Martin et al., 1989; Mullennix & Pisoni, 1990; Mullennix et al., 1989; Palmeri et al., 1993). However, we were not testing whether speech perception was facilitated by matching, only if sensitivity was affected by matching. It is entirely possible that we would find evidence of talker-sensitivity using a different task. Some researchers have argued that learned indexical information is accessed at a relatively higher level of speech processing (Mattys & Liss, 2008; McLennan, 2006; McLennan & Luce, 2005). Since the AX discrimination test taps into a relatively lower level of speech processing, that is, at the auditory rather than phonetic level (e.g., Pisoni, 1973), participants could have processed test stimuli before they accessed learned indexical information, and thus not shown talker-specificity effects.

How did the participants in the offset condition acquire the overarching contrast from talker-specific distributions seemingly without acquiring separate talker-specific categories? One possibility is that after learning that each individual talker produced a bimodal distribution, participants noted that there was something consistent across these distributions, namely, the

distance between frequency peaks (i.e., 100 ms between frequency peaks). Since the AX discrimination test assessed participants' sensitivity to an acoustic discrepancy between two test stimuli, knowing that segments that differ from each other by 100 ms belong to different categories could have facilitated their performance in the AX discrimination regardless of whether the test stimuli were the canonical exemplars of the target category for the talker who produce that specific test item.

Summary

In this paper, we examined whether adults can learn two phonetic categories from in input in which talker-dependent variation in the acoustic signal introduces potential ambiguity in the categorization of speech sounds in input. The results of our experiment suggest that adults can overcome the potential ambiguity, possibly by using indexical information to track statistical information in the input separately by talker. This suggests that distributional learning can account for the learning of phonetic categories from natural language input even in the more realistic case where talker-dependent variation occurs. This study demonstrates the robustness and the complexity of statistical language learning.

References

- Bion, R. A., Miyazawa, K., Kikuchi, H., & Mazuka, R. (2013). Learning phonemic vowel length from naturalistic recordings of Japanese infant-directed speech. *PloS One*, 8(2), e51594. doi:10.1371/journal.pone.0051594
- Boersma, P. (2002). Praat, a system for doing phonetics by computer. *Glott International*, 5(9/10), 341–345.
- Bradlow, A. R., Torretta, G. M., & Pisoni, D. B. (1996). Intelligibility of normal speech I: Global and fine-grained acoustic-phonetic talker characteristics. *Speech Communication*, 20(3), 255-272.
- Craik, F. I. & Kirsner, K. (1974). The effect of speaker's voice on word recognition. *The Quarterly Journal of Experimental Psychology*, 26(2), 274–284.
- Creelman, C. D. (1957). Case of the unknown talker. *Journal of the Acoustical Society of America*, 25, 655.
- Cumming, R. (2011). The effect of dynamic fundamental frequency on the perception of duration. *Journal of Phonetics*, 39(3), 375-387.
- Goldinger, S. D. (1996). Words and voices: episodic traces in spoken word identification and recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(5), 1166–1183.
- Goldinger, S. D., Pisoni, D. B., & Logan, J. S. (1991). On the nature of talker variability effects on recall of spoken word lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17(1), 152-162.
- Hagiwara, R. (1997). Dialect variation and formant frequency: The American English vowels revisited. *The Journal of the Acoustical Society of America*, 102(1), 655-658.

- Hillenbrand, J., Getty, L. A., Clark, M. J., & Wheeler, K. (1995). Acoustic characteristics of American English vowels. *The Journal of the Acoustical society of America*, 97(5), 3099–3111.
- Hirata, Y. (2004). Training native English speakers to perceive Japanese length contrasts in word versus sentence contexts. *The Journal of the Acoustical Society of America*, 116(4), 2384–2394.
- Hirata, Y., Whitehurst, E., & Cullings, E. (2007). Training native English speakers to identify Japanese vowel length contrast with sentences at varied speaking rates. *The Journal of the Acoustical Society of America*, 121(6), 3837–3845.
- Johnson, K. (1997). Speech perception without speaker normalization: An exemplar model. In K. Johnson & J. W. Mullennix (Eds.), *Talker variability in speech processing*, (pp. 145–165). New York: Academic Press.
- Johnson, K. & Mullennix, J. W. (Eds.). (1997). *Talker variability in speech processing*. New York: Academic Press.
- Jongman, A., Wayland, R., & Wong, S. (2000). Acoustic characteristics of English fricatives. *The Journal of the Acoustical Society of America*, 108(3), 1252-1263.
- Labrune, L. (2012). *The Phonology of Japanese*. Oxford: Oxford University Press.
- Lindblom, B. E. & Studdert-Kennedy, M. (1967). On the role of formant transitions in vowel recognition. *The Journal of the Acoustical Society of America*, 42(4), 830–843.
- Logan, J. S., Lively, S. E., & Pisoni, D. B. (1991). Training Japanese listeners to identify English /r/ and /l/: a first report. *The Journal of the Acoustical Society of America*, 89(2), 874–886.

- Macmillan, N. A. & Creelman, C. D. (2005). *Detection theory: a user's guide*. Mahwah: Lawrence Erlbaum Associates, Inc.
- Mann, V. A. (1980). Influence of preceding liquid on stop-consonant perception. *Perception & Psychophysics*, 28(5):407–412.
- Mann, V. A. & Repp, B. H. (1980). Influence of vocalic context on perception of the [f]-[s] distinction. *Perception & Psychophysics*, 28(3), 213–228.
- Martin, C. S., Mullennix, J. W., Pisoni, D. B., & Summers, W. V. (1989). Effects of talker variability on recall of spoken word lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15(4), 676.
- Mattys, S. L. & Liss, J. M. (2008). On building models of spoken-word recognition: when there is as much to learn from natural “oddities” as artificial normality. *Perception & Psychophysics*, 70(7), 1235–1242.
- Maye, J. (2000). *Learning speech sound categories from statistical information* (Doctoral dissertation, University of Arizona).
- Maye, J., Werker, J. F., & Gerken, L. (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition*, 82(3), B101–B111.
- McLennan, C. T. (2006). The time course of variability effects in the perception of spoken language: changes across the lifespan. *Language and Speech*, 49(1), 113–125.
- McLennan, C. T. & Luce, P. A. (2005). Examining the time course of indexical specificity effects in spoken word recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(2), 306–321.
- Mullennix, J. W. & Pisoni, D. B. (1990). Stimulus variability and processing dependencies in speech perception. *Perception & Psychophysics*, 47(4), 379–390.

- Mullennix, J. W., Pisoni, D. B., & Martin, C. S. (1989). Some effects of talker variability on spoken word recognition. *The Journal of the Acoustical Society of America*, 85(1), 365–378.
- Pajak, B. (2012). *Inductive inference in non-native speech processing and learning* (Doctoral dissertation, University of California, San Diego).
- Palmeri, T. J., Goldinger, S. D., & Pisoni, D. B. (1993). Episodic encoding of voice attributes and recognition memory for spoken words. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19(2), 309.
- Pisoni, D. B. (1973). Auditory and phonetic memory codes in the discrimination of consonants and vowels. *Perception & Psychophysics*, 13(2), 253–260.
- Repp, B. H. & Mann, V. A. (1981). Perceptual assessment of fricative–stop coarticulation. *The Journal of the Acoustical Society of America*, 69(4), 1154–1163.
- Repp, B. H. (1981). Two strategies in fricative discrimination. *Perception & Psychophysics*, 30(3), 217–227.
- Summerfield, Q. (1975). How a full account of segmental perception depends on prosody and vice versa. In A. Cohen & S. G. Nooteboom (Eds.), *Structure and process in speech perception*, (pp. 51–68), Berlin: Springer-Verlag.