

Lexically-Mediated Compensation for Coarticulation in Older Adults

Sahil Luthra (sahil.luthra@uconn.edu)

Department of Psychological Sciences, University of Connecticut, Storrs, CT, 06269 USA

Giovanni Peraza-Santiago (giovanni.peraza-santiago@uconn.edu)

Department of Psychological Sciences, University of Connecticut, Storrs, CT, 06269 USA

David Saltzman (david.saltzman@uconn.edu)

Department of Psychological Sciences, University of Connecticut, Storrs, CT, 06269 USA

Anne Marie Crinnion (anne.crinion@uconn.edu)

Department of Psychological Sciences, University of Connecticut, Storrs, CT, 06269 USA

James S. Magnuson (james.magnuson@uconn.edu)

Department of Psychological Sciences, University of Connecticut, Storrs, CT, 06269 USA
Basque Center on Cognition, Brain and Language, Mikeletegi Pasealekua 69, Donostia-San Sebastián, Spain
Ikerbasque. Basque Foundation for Science, Bilbao, Spain

Abstract

The claim that contextual knowledge exerts a top-down influence on sensory processing is supported by evidence for lexically-mediated compensation for coarticulation (LCfC) in spoken language processing. In this phenomenon, a lexically restored context phoneme (e.g., the final phoneme in *Christma#* or *fooli#*) influences perception of a subsequent target phoneme (e.g., a phoneme ambiguous between /t/ and /k/). A recent report shows that carefully vetted materials produce robust, replicable LCFc effects in younger adults (18-34 years old). Here, we asked whether we would observe LCFc in a sample of older adults (aged 60+). This is of interest because older adults must often contend with age-related declines in sensory processing, with previous research suggesting that older adults may compensate for age-related changes by relying more strongly on contextual knowledge. We observed robust LCFc effects in younger and older samples, with no significant difference in the effect size between age groups.

Keywords: spoken word recognition; interactive models; activation feedback; aging; speech perception

Introduction

In order to recognize a spoken word, listeners must map the incoming auditory signal onto a known mental representation. Thus, spoken word recognition necessarily depends on a series of bottom-up transformations applied to the signal. At the same time, studies have provided evidence that high-level contextual information, such as lexical knowledge, can guide a listener's interpretation of the speech signal, particularly when the auditory input is ambiguous (e.g., Ganong, 1980). For instance, a listener's interpretation of a speech sound that is ambiguous between /s/ and /ʃ/ ("sh") is guided by the context in which it appears, with listeners more likely to characterize the ambiguous sound as /s/ if it is preceded by *Christma_* and as /ʃ/ if it is preceded by *fooli_*. As such, spoken word recognition may be appropriately characterized as a balancing act, representing the interplay

between ascending auditory processes and descending influences of higher-level cognitive processes.

While the importance of these contextual influences is well-established, there has been a long-standing debate over how these influences manifest computationally (e.g., McClelland, Mirman, & Holt, 2006; McQueen, Norris, & Cutler, 2006). Proponents of interactive models, such as the TRACE model of speech perception (McClelland & Elman, 1986), propose that when low-level auditory information propagates forward, relevant lexical representations are activated; critically, lexical representations then feed information *backward* to help boost activation of corresponding phonemes. Such models suggest that higher levels of lexical processing can *directly* modulate low-level sensory processing, a feature that is particularly advantageous during noisy conditions because it permits contextual information to constrain interpretation of the speech signal to relevant candidates (Magnuson, Mirman, Luthra, Strauss, & Harris, 2018). Critics of interactive models argue that lexical feedback would impede veridical perception of the auditory signal, leading the model to perceive "only its own predictions" and therefore causing rampant hallucinations (Norris, McQueen & Cutler, 2016).

In a seminal study, Elman and McClelland (1988) argued in favor of interactive models based on evidence for *lexically-mediated compensation for coarticulation* (LCfC). This paradigm couples two phenomena in order to isolate the locus of context effects. The first phenomenon is *compensation for coarticulation*. When a talker has to produce a speech sound that is produced relatively far back in the mouth (e.g., /ʃ/, /g/, /k/, /r/) immediately after a speech sound produced near the front (e.g., /s/, /d/, /t/, /l/), the talker may not reach the intended place of articulation (PoA) as a result of motoric constraints of speech production (i.e., the fact that the articulatory gestures for successive speech sounds overlap in time and so are *coarticulated*). As a result, the second speech

sound might be produced with an ambiguous PoA. (Note that the analogous context effect is observed when a more back PoA precedes a more front PoA). Researchers had long established that listeners are sensitive to these kinds of contingencies and compensate for this coarticulatory effect. For instance, if listeners heard a sound with an ambiguous PoA (e.g., a sound ambiguous between /t/ and /k/), but this sound was preceded by a sound with a clear front PoA (e.g., /s/), listeners would perceive the ambiguous sound as having a back PoA (e.g., Mann & Repp, 1981; Repp & Mann, 1981, 1982). Though some researchers have suggested that this effect is driven by general auditory processes and not motoric knowledge (e.g., Holt & Lotto, 2008), compensation for coarticulation is critically agreed to have a sub-lexical locus.

The second phenomenon is *phoneme restoration*, a context effect where listeners appear to fill in a sound that has been replaced by noise (Samuel, 1981, 2001; Warren, 1970). Elman and McClelland (1988) reasoned that if lexical information could directly modulate a sub-lexical process, as predicted by an interactive model, then a lexically-restored phoneme should be able to drive compensation for coarticulation. They merged the two paradigms by pairing a context word with an ambiguous final phoneme that could only resolve to one option (e.g., *Christma#* or *fooli#*) with a target word where the initial phoneme varied along a front-back PoA continuum (e.g., *#apes*, with # taken from a /t/-/k/ continuum). Critically, the direction of the compensation for coarticulation effect depended on whether lexical information restored the ambiguous context phoneme as having front PoA (e.g., /s/) or back PoA (e.g., /l/).

While the logic of the LCfC paradigm represents a gold standard for determining whether context has a top-down influence on sensory processing, the field has been stymied by inconsistent results in subsequent LCfC studies (Pitt & McQueen 1998; Magnuson et al., 2003; McQueen, Jesse & Norris 2009), making it difficult for the field to adjudicate between models with and without feedback.

Recently, Luthra et al. (2021) speculated that the weak results in previous LCfC studies may have been driven in part by characteristics of the stimuli. In particular, they noted that most previous studies had not pretested stimuli to ensure that they would elicit the requisite phoneme restoration effects (e.g., that lexical knowledge would guide a listener to interpret the final segment of *Christma#* as /s/) and that they would elicit classic compensation for coarticulation effects (e.g., that an unambiguous context stimulus like *Christmas* would induce the expected compensation for coarticulation effects on target items from a subsequent front-back PoA continuum) *before* combining items in the LCfC paradigm. Luthra et al. extensively piloted potential context and target stimuli to ensure that they would elicit these baseline effects. To assess potential context stimuli, Luthra et al. presented one group of listeners with items on a word-nonword continuum (e.g., *abolish-aboliss**) and one group of listeners with stimuli that were trimmed to yield a nonword-nonword continuum (e.g., *ish-iss**); if items showed a lexical effect (i.e., at some step, participants made more lexically

consistent responses for the intact stimuli than for the trimmed stimuli), they were submitted to additional piloting. To assess potential target stimuli, the authors tested whether compensation for coarticulation was observed when targets were preceded by context items with unambiguous endings (e.g., *abolish #apes*). After stimuli were pretested, Luthra et al. conducted a well-powered, pre-registered study to test for LCfC; a post-hoc power analysis indicates that Luthra et al. had power of 0.99 for detecting the effect of interest. They observed robust LCfC effects in a sample of young adults and also replicated these findings in a second independent sample, providing strong evidence for modulation of sublexical processing via lexical feedback.

In the current work, we aim to extend that previous study by examining how LCfC effects might change across the lifespan. A consideration of aging can allow for a richer understanding of the interactions between top-down and bottom-up processing, as older adults often endure age-related changes in hearing abilities and in cognitive processing (Rogers & Peelle, submitted). While overall processing tends to be slower in older adults than in their younger peers (Salthouse, 1996), crystallized cognitive abilities, such as lexical knowledge, appear to be relatively preserved across the lifespan, and as such, older adults may rely more on high-level context to compensate for age-related declines in hearing (Mattys & Scharenborg, 2014; Pichora-Fuller, 2008; Rogers, Jacoby, & Sommers, 2012). For the current study, we hypothesized that because older adults rely more heavily on lexical knowledge when processing ambiguous phonetic information, LCfC effects might be larger in a sample of older adults than in a sample of younger adults.

Experiment 1

In Experiment 1, we tested whether LCfC would be observed in a sample of older adults and how the size of the effect might compare to a sample of younger adults. Data for the younger sample were taken from the first experiment of the recent study reviewed above (Luthra et al., 2021), available from an Open Science Framework repository (<https://osf.io/q8c3z/>). Data for the older adults sample were collected and analyzed in accordance with a plan we pre-registered on the Open Science Framework (<https://osf.io/4zvyu/>).

Method

Materials Stimuli were taken directly from Luthra et al. (2021). There were 4 context items (for lexically-based phoneme restoration) with ambiguous final phonemes (*isolate*/isolake*, *maniac*/maniat*, *pocketful*/pocketfur* and *questionnaire*/questionnail*) as well as 5 target continua for compensation for coarticulation (*same/shame*, *sell/shell*, *sign/shine*, *sip/ship* and *sort/short*).

As described by Luthra et al. (2021), materials were prepared by first recording a native speaker of American English, who produced clear versions of the endpoints for each target continuum (e.g., *same* and *shame*). 11-step

continua were constructed in STRAIGHT (Kawahara et al., 2008); STRAIGHT requires the experimenter to identify temporal and spectral landmarks in each of the endpoint stimuli and then interpolates between the endpoints in equal steps to create the desired continuum. For each target continuum, the most ambiguous step (as identified during pilot testing) was identified as step 0 and all other steps were expressed relative to step 0. In Experiment 1 of the current study, listeners heard five steps (-2, -1, 0, 1 and 2) from each target continuum, following the approach of Luthra et al.

Procedure The current experiment was programmed using the online experiment builder Gorilla (Anywl-Irvine et al., 2020) and data were collected through the Prolific platform (<https://prolific.sc>). After providing informed consent, participants answered a series of demographics questions and completed a stereo listening test to ensure they were using headphones (Woods et al., 2017).

In the experiment proper, each trial consisted of an ambiguous context word (e.g., *isola#*) immediately followed by the critical target word. Participants had to decide whether the target word began with a front or back place-of-articulation (with /s/ or /ʃ/, respectively) by pressing the relevant key. Subjects completed 6 blocks of this task, with each block consisting of all 100 possible trials (4 context items x 5 target continua x 5 steps / target continuum) in random order. Response mappings (i.e., whether the ‘s’ button response was on the left or on the right) were counterbalanced across participants. Trials timed out after 6 seconds, and there was a 1-second inter-trial interval. In total, the experiment took approximately 45 minutes to complete, and participants were paid \$9 for their participation, consistent with Connecticut’s minimum wage (\$12/hour) at the time of the study.

Participants 69 older adults were recruited for the current experiment. These participants self-identified as monolingual speakers of English above the age of 65 who lived in the United States. All subjects reported having normal or corrected-to-normal vision, no hearing difficulties, and no language-related disorders, and all subjects had at least a 90% approval rating on Prolific (indicating a high rate of compliance in previous studies).

As specified in our pre-registration, we excluded participants who failed the headphone screening test twice as well as participants with poor behavioral performance (subjects with <80% classification accuracy for clear endpoints from target continua and/or who failed to respond to 10% or more of trials). This resulted in a sample of 40 older adults (25 female, 15 male; mean age: 69, range: 65-75).

We compared the performance of the older adults to an archival sample of 40 younger adults (22 female, 18 male; mean age: 27, range: 18-34) taken from Luthra et al. (2021). Notably, a power analysis using the effect size from this archival sample suggested that we only needed 15 participants to estimate the effect of the context item (i.e., the LCfC effect) with power of 0.90 at an alpha level of 0.05. However, it is unclear whether this is an appropriate estimate for the older adult sample, as LCfC effects might differ in size

across age ranges. By analyzing data from 40 older adult participants, we ensured that our samples were matched in size. Furthermore, by using a substantially larger sample size than indicated by the power analysis, we are better able to test whether there is an interaction between age group and the size of the LCfC effect.




Figure 1: Comparably robust lexically-mediated compensation for coarticulation (LCfC) effects in samples of older adults (left) and younger adults (right; archival data from Luthra et al., 2021). Participants heard target items taken from a front-back place-of-articulation (PoA) continuum; the x-axis shows the continuum step, with 0 representing the most ambiguous step. The y-axis indicates how often subjects responded that the first segment of the target stimulus following the context item had a front PoA. Solid lines represent the average response when a preceding context item ended with an implied back PoA (e.g., *questionnai#*) and dotted lines represent a context item that implied a front PoA for the final segment (e.g., *isola#*). As expected, responses shifted depending on the lexically implied PoA for the final segment of the context item.

Results

We observed strong LCfC effects in our sample of older adults, as shown in Figure 1.

Older adult data were submitted to a mixed effects logistic regression implemented in R (R Core Team, 2019) that tested for fixed effects of Context Bias (whether lexical information in the context item implied a front or back PoA for the final segment; front/back, coded with a [1,-1] contrast) and Step (scaled). The model also included random by-subject and by-item intercepts as well as random by-subject by-item slopes for Context Bias and Step; note that the items here correspond to the target items (*same/shame*, *sign/shine*, etc.). This random effect structure is both the maximal one (Barr et al., 2013) and the most parsimonious one (Matuscheck et al., 2017), as simplifications to the random effect structure led to a poorer model fit. All regressions were implemented using the *mixed* function in the “afex” package (Singmann et al., 2020); this function interfaces with the *glmer* function in the “lme4” package (Bates et al., 2015) and uses likelihood ratio tests to evaluate the significance of all fixed factors and their

interactions. Specifically, the model compares how well the data are captured by a model that includes the fixed effect of interest relative to a model that does not include the effect. Thus, significant effects are those that improve model fit, as indicated by a significant chi-square test statistic.

We observed a significant effect of Context Bias in our older adult sample, $\chi^2(1) = 14.16$, $p < 0.0001$, indicating LCfC in this group. We also observed a significant effect of Step, $\chi^2(1) = 21.69$, $p < 0.0001$, indicating that the proportion of front-PoA responses was lower for target stimuli that had a relatively back PoA. There was no significant interaction between these factors, $p = 0.60$.

To test for possible influences of age on the size of the LCfC effect, we compared our older adults sample to the archival sample of younger adults. Specifically, we conducted a regression with fixed factors of Context Bias, Step, and Age Group (older/younger, coded with a [1,-1] contrast), random by-subject intercepts, random by-subject slopes and interactions for Context Bias and Step, random by-item intercepts, and random by-item slopes and interactions for Context Bias, Step and Age Group. As before, the maximal model was also the most parsimonious. We observed significant effects of Context Bias, $\chi^2(1) = 13.62$, $p = 0.0002$, and of Step, $\chi^2(1) = 21.16$, $p < 0.0001$. We also observed a significant interaction between Step and Age Group, $\chi^2(1) = 16.29$, $p < 0.0001$, though this effect was not of theoretical interest. No other effects were significant.

It is particularly striking that there was no difference in the size of the LCfC effect across age groups (i.e., no interaction between Context Bias and Age Group). To quantify the size of the LCfC effect in each group, we calculated the difference between how often subjects made /s/ responses when the context item ended with an implied back PoA compared to when it ended with an implied front PoA. For older adults, the effect size was 7.8%, and for younger adults, the effect size was 7.1%.

Finally, we compared the response times for older and younger adults; note that this analysis was not included in our pre-registration for the study and as such constitutes an exploratory analysis. To eliminate outlier response times, we only considered responses less than 2000 ms (97.5% of the data). We then conducted a linear mixed effects regression with fixed factors of Context Bias, Step, and Age Group. As before, the random effect structure was selected through a backward-stepping procedure, starting with the maximal random effects structure. Through this process, we selected a random effects structure with (1) random by-subject intercepts, (2) random by-subject slopes and interactions for Context Bias and Step, (3) random by-item intercepts, and (4) random by-item slopes and interactions for Context Bias, Step and Age Group. However, in contrast to previous models, our model did not include random correlations between random slopes and intercepts (see Barr et al., 2013 for discussion). We specified in the model that the data were linked to a Gamma distribution with an identity function, following the recommendation of Lo and Andrews (2015). Note that because some response times were measured from

stimulus offset, some data were close to the negative boundary, which can produce errors during model fitting, since the Gamma distribution excludes negative values. We thus added a constant value of 50 ms to all response times for the purposes of this analysis; these shifted reaction times were used as the dependent variable. Note that the addition of a constant value only influences the intercept of the model and does not influence the estimates for the fixed effects. Of interest was whether we would observe any differences in response time as a function of Age Group. Though the mean response time for older adults (455 ms, untransformed) was longer than the mean response time for younger adults (418 ms), we did not observe any effects of Age Group in this analysis. The only significant effects were of Context Bias, $\chi^2(1) = 9.48$, $p = 0.002$, and of Step, $\chi^2(1) = 5.60$, $p = 0.02$. Because neither effect is of theoretical interest, we do not discuss them further.

Discussion

In Experiment 1, we observed LCfC effects in a sample of older adults; to our knowledge, this is the first demonstration of LCfC in this age group. Strikingly, the size of the LCfC effect was comparable between our sample of older adults and an archival sample of younger adults collected for a previous study (Luthra et al., 2021), as evidenced by the lack of interactions between Context Bias and Age Group as well as the comparable effect sizes between groups.

In examining the data in Figure 1, however, it seemed possible to us that older adults may have had a slightly more pronounced LCfC effect than younger adults at more intermediate continuum steps (i.e., step 0), but that our ability to detect this might be dampened by performance on less central steps (i.e., steps -2 and 2). We therefore decided to conduct a follow-up experiment that would look more closely at the intermediate range of the target continuum.

Experiment 2

In Experiment 2, we repeated Experiment 1 but selected target stimuli from a more limited range around the most ambiguous continuum step. In this way, we aimed to improve our power to detect differences in the middle of the target continuum, where LCfC effects are likely to be most pronounced. Note that Experiment 2 was not pre-registered.

Methods

Materials For Experiment 2, we created 21-step target continua using the same parameters as in the previous study by Luthra et al. (2021). As illustrated in Figure 2, this allowed us to select steps that spanned a more narrow region around the most ambiguous step. For instance, Experiment 1 used steps 3-7 from the *sort/short* continuum since step 5 was identified in pilot testing as having the most ambiguous PoA; Luthra et al. re-expressed this as steps -2 to 2 so that 0 was the most ambiguous step. As shown in Figure 3, step 5 on the 11-step continuum corresponds to step 9 on the 21-step continuum. Thus, for the *sort/short* continuum in Experiment 2, we selected steps 7 through 11 from the 21-step continuum,

which spans the same range as steps 4-6 from the 11-step continuum. We then relabeled our steps so that 0 represented the most ambiguous step. Thus, while Experiment 1 used steps -2 (front), -1, 0 (ambiguous), 1 and 2 (back), Experiment 2 used steps -1 (front), -0.5, 0 (ambiguous), 0.5, and 1 (back).




Figure 2: To construct new stimuli for Experiment 2, we created 21-step continua instead of the 11-step continua used in Experiment 1 and by Luthra et al. (2021).

Procedure The procedure was the same as in Experiment 1. **Participants** For Experiment 2, we recruited 50 young adults and 56 older adults who had not participated in Experiment 1. As before, we excluded participants who failed the headphone screening test twice, and we also excluded participants who failed to respond to 10% or more of the trials. In contrast to Experiment 1, we did not exclude participants with low accuracy in their classification of the target continuum endpoints, as participants in Experiment 2 did not hear unambiguous endpoints from the target continua. One additional participant was randomly excluded to equalize the number of participants in each age group. Consequently, data from 40 young adults (25 female, 15 male; mean age: 30, range: 19-34) and 40 older adults (26 female, 14 male; mean age: 68, range: 60-76) were included in analyses for Experiment 2.

Note that the age range for the older adults in Experiment 2 (60-76) is slightly larger than the range in Experiment 1 (65-75). For Experiment 1, our pre-registered recruitment strategy involved recruiting individuals above 65, though we noted that if we had difficulty with recruitment, we would reduce the lower bound to 60. Though Experiment 2 was not pre-registered, we opted to follow the same recruitment strategy as pre-registered for Experiment 1. Thus, when we encountered difficulty recruiting a full sample of older adults aged 65+ in Experiment 2, we opted to reduce the lower bound to 60.

Results

Results from Experiment 2 are shown in Figure 3. As before, we observed robust LCfC effects, as participants' interpretation of an ambiguous target phoneme depended on whether the context item had a lexically implied front or back PoA for its final segment (e.g., *isola#* or *questionnai#*).

To assess the size of the LCfC effects across groups, we conducted a regression analysis following the same approach

as in Experiment 1. As before, we observed significant effects of Context Bias (front/back), $\chi^2(1) = 9.41, p = 0.002$, and of Step, $\chi^2(1) = 14.88, p = 0.0001$. No other effects were significant. Effect sizes were computed as before; we observed a numerically larger effect in our older adult sample (6.1%) than in our younger sample (4.3%).




Figure 3: In Experiment 2, we tested for LCfC effects using target continua with a more limited range (half-step increments from -1 to 1 instead of full-step increments from -2 to 2, as in Experiment 1). We observed robust LCfC effects in both age groups, with no difference in effect size as a function of age.

Finally, we tested for potential differences in processing speed between groups by submitting response times to a linear mixed effects regression analysis following the same approach as above. Note that as before, we only included response times faster than 2000 ms (96.6% of the data). Though responses were numerically faster for the younger subjects (mean: 462 ms) than for older subjects (mean: 485 ms), we did not observe any significant effects of Age Group. We did observe significant effects of Context Bias, $\chi^2(1) = 3.92, p = 0.05$, and of Step, $\chi^2(1) = 6.63, p = 0.01$, as was also seen in Experiment 1.

Discussion

In Experiment 2, we replicated the core findings from Experiment 1. In particular, both older adults and younger adults exhibited robust LCfC effects, but there was no significant difference in the effect size as a function of age. We note that overall, responses in Experiment 2 were relatively close to chance across the entire range of the target continua. We suspect that this is because participants in Experiment 2 never heard clear endpoints for the target stimuli (i.e., they never heard a clear production of *same* or of *shame*), and so they may have been more uncertain about how to categorize the stimuli in general. It is striking that LCfC effects were still observed in both age groups despite this increased uncertainty.

General Discussion

Across two experiments, we observed robust LCfC, in that a lexically implied PoA (e.g., the front PoA implied by *isola*) influenced perception of a subsequent target segment with an ambiguous PoA. Critically, if we accept the assumptions of Pitt and McQueen (1998), when transitional probabilities are not confounded with lexical status (Luthra et al., 2021, show that there is no single transitional probability that can explain all the positive effects), LCfC can only be explained by computational models that allow for lexical-to-sublexical feedback. In such interactive models, lexical information modulates the way that the context phoneme influences a sublexical process (compensation for coarticulation). In autonomous models, there is no way for lexical knowledge to influence sublexical processing; it can only influence sublexical decisions on such accounts. Thus, our findings provide strong evidence in favor of interactive accounts of spoken word recognition.

In the present study, we also demonstrated that LCfC effects can be observed in older adults, obtaining this result in two separate samples; to our knowledge, previous studies have not shown LCfC effects in older adults. Thus, our data support the idea that interactions between top-down and bottom-up processing are present across the lifespan.

Notably, the LCfC effects in our older and younger adult samples were of approximately equal magnitude. We had hypothesized that because older adults reportedly rely relatively heavily on higher-level knowledge (e.g., Rogers et al., 2012), LCfC effects might be larger in older adults. However, the present results may still be consistent with an account in which older adults are upweighting high-level knowledge compared to their younger peers. Age-related declines in sensory processing are often observed in older adults (Rogers & Peelle, submitted), and previous work has suggested that older adults may mitigate declines in sensory processing by relying more strongly on contextual knowledge (c.f., Mattys & Scharenborg, 2014; Pichora-Fuller, 2008). It is possible that older adults in the current study faced age-related declines in sensory processing, even if response times were comparable between older and younger adults. If this were the case, older adults may have needed to rely more on lexical knowledge in order to show LCfC effects of comparable size to the younger adults. Assessing this hypothesis will require additional studies with finer-grained measures of sensory processing (e.g., performance on a [non-lexical] compensation for coarticulation task [e.g., Mann & Repp, 1981] and/or comprehensive assessments of individual differences.

Additionally, it is possible that the lack of age-related differences in the size of the LCfC effect is due to a potential ceiling effect. Both young adults and older adults have robust lexical knowledge, and as such, both groups might be expected to show strong LCfC effects. Future work might assess the size of LCfC effects in groups with less lexical knowledge, such as children.

Finally, it is possible that the perceptual processes listeners bring to bear in the current task interact minimally with

broader aspects of cognition. This would be consistent with some previous work, such as results from Zhang and Samuel (2014), who found no effect of cognitive load on lexically-guided perceptual learning. Thus there may be little reason to expect normal hearing older adults to perform differently from younger adults.

In closing, this study demonstrates a robust lexically mediated compensation for coarticulation effect in older adults; to our knowledge, this has not been shown in previous studies. Thus, the present results provide an important contribution to the literature characterizing the interplay between top-down and bottom-up processing in older adults.

Acknowledgments

This research was supported by NSF BCS-PAC 1754284 and NSF NRT 1747486 (PI: JSM). This research was also supported in part by the Basque Government through the BERC 2018-2021 program, and by the Agencia Estatal de Investigación through BCBL Severo Ochoa excellence accreditation SEV-2015-0490. SL was supported by an NSF Graduate Research Fellowship, and AMC was supported by NIH T32 DC017703. Data and analysis code are publicly available at <https://osf.io/q5392/>.

References

- Anywl-Irvine, A. L., Massonnié, J., Flitton, A., Kirkham, N., & Evershed, J. K. (2020). Gorilla in our midst: An online behavioral experiment builder. *Behavior Research Methods*, 52, 388-407.
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255-278.
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1-48.
- Ganong, W. F. (1980). Phonetic categorization in auditory perception. *Journal of Experimental Psychology: Human Perception and Performance*, 6(1), 110-125.
- Holt, L. L., & Lotto, A. J. (2008). Speech perception within an auditory cognitive science framework. *Current Directions in Psychological Science*, 17(1), 42-46.
- Kawahara, H., Morise, M., Takahashi, T., Nisimura, R., Irino, T., & Banno, H. (2008). Tandem-STRAIGHT: A temporally stable power spectral representation for periodic signals and applications to interference-free spectrum, F0, and aperiodicity estimation. In *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing* (pp. 3933-3936).
- Lo, S., & Andrews, S. (2015). To transform or not to transform: Using generalized linear mixed models to analyse reaction time data. *Frontiers in Psychology*, 6, 1-16.
- Luthra, S., Peraza-Santiago, G., Beeson, K., Saltzman, D., Crinnion, A. M., & Magnuson, J. S. (2021). Robust lexically-mediated compensation for coarticulation:

- Christmash time is here again. *Cognitive Science*, 45(4), 1-20.
- Magnuson, J. S., McMurray, B., Tanenhaus, M. K., & Aslin, R. N. (2003). Lexical effects on compensation for coarticulation: The ghost of Christmash past. *Cognitive Science*, 27(2), 285–298.
- Magnuson, J. S., Mirman, D., Luthra, S., Strauss, T., & Harris, H. D. (2018). Interaction in spoken word recognition: Feedback helps. *Frontiers in Psychology*, 9.
- Mann, V. A., & Repp, B. H. (1981). Influence of preceding fricative on stop consonant perception. *Journal of the Acoustical Society of America*, 69(2), 548–558.
- Mattys, S. L., & Scharenborg, O. (2014). Phoneme categorization and discrimination in younger and older adults: A comparative analysis of perceptual, lexical, and attentional factors. *Psychology and Aging*, 29(1), 150-162.
- Matuschek, H., Kliegl, R., Vasishth, S., Baayen, H., & Bates, D. (2017). Balancing Type I error and power in linear mixed models. *Journal of Memory and Language*, 94, 305–315.
- McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, 18(1), 1-86.
- McClelland, J. L., Mirman, D., & Holt, L. L. (2006). Are there interactive processes in speech perception? *Trends in Cognitive Sciences*, 10(8), 363-369.
- McQueen, J. M., Jesse, A., & Norris, D. (2009). No lexical-prelexical feedback during speech perception or: Is it time to stop playing those Christmas tapes? *Journal of Memory and Language*, 61(1), 1–18.
- McQueen, J. M., Norris, D., & Cutler, A. (2006). Are there really interactive processes in speech perception? *Trends in Cognitive Sciences*, 10(12), 533.
- Norris, D., McQueen, J. M., & Cutler, A. (2016). Prediction, Bayesian inference and feedback in speech recognition. *Language, Cognition and Neuroscience*, 31(1), 4-18.
- Pichora-Fuller, M. K. (2008). Use of supportive context by younger and older adult listeners: Balancing bottom-up and top-down information processing. *International Journal of Audiology*, 47, S72-S82.
- Pitt, M. A., & McQueen, J. M. (1998). Is compensation for coarticulation mediated by the lexicon? *Journal of Memory and Language*, 39, 347–370.
- R Core Team. (2019). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org/>.
- Rogers, C. S., Jacoby, L. L., & Sommers, M. S. (2012). Frequent false hearing by older adults: The role of age differences in metacognition. *Psychology and Aging*, 27(1), 33-45.
- Rogers, C. S., & Peelle, J. E. Interactions between audition and cognition in hearing loss and aging. In L. L. Holt & A. J. Lotto (Eds.), *The Auditory Cognitive Neuroscience of Speech Perception*. New York: Springer. Submitted.
- Salthouse, T. A. (1996). The processing-speed theory of adult age differences in cognition. *Psychological Review*, 103(3), 403-428.
- Samuel, A. G. (1981). Phonemic restoration: Insights from a new methodology. *Journal of Experimental Psychology: General*, 110, 474-494.
- Samuel, A. G. (2001). Knowing a word affects the fundamental perception of the sounds within it. *Psychological Science*, 12, 348-351.
- Singmann, H., Bolker, B., Westfall, J., & Aust, F. (2020). afex: Analysis of Factorial Experiments. R package version 0.26. <https://CRAN.R-project.org/package=afex>.
- Warren, R. M. (1970). Perceptual restoration of missing speech sounds. *Science*, 167(3917), 392–393.
- Woods, K. J. P., Siegel, M. H., Traer, J., & McDermott, J. H. (2017). Headphone screening to facilitate web-based auditory experiments. *Attention, Perception, and Psychophysics*, 79(7), 2064–2072.
- Zhang, X., & Samuel, A. G. (2014). Perceptual learning of speech under optimal and adverse conditions. *Journal of Experimental Psychology: Human Perception and Performance*, 40(1), 200–217.