Perception of Coda Nasal Consonants in English: Effects of Linguistic Experience

Megan L. Risdal

1 Introduction

It has been noted among sociolinguistics that speakers of African American (Vernacular) English (AAVE) often exhibit a propensity to delete (or more likely reduce) nasal consonant articulations and retain the preceding nasalized vowel (Wolfram 1989). For example, a speaker of this dialect might say [gẽ] or [gẽʔ] instead of the unreduced [gem] or [gẽm] for the word "game." Unlike coda t/d deletion, coda nasal consonant deletion has not been studied extensively in AAVE. Descriptions of coda nasal consonant deletion among speakers of AAVE have been limited to anecdotal evidence among sociolinguists and small-scale analysis of interview speech (Risdal 2015). The dearth of production studies focused on this question is likely due in part a result of the *observer's paradox* and the prevalence of bidialectalism among speakers of African American English (Labov 1972:pg. 209). In light of these limitations and their underlying reasons, experimental study of the perception of the effects of coda nasal consonant deletion among in- and out-groups presents an attractive alternative to investigating this phenomenon. In this paper, I propose an experimental method with the ultimate goal of providing an account of coda nasal consonant deletion based on the perceived similarity between the nasal consonant and the preceding nasalized vowel (Steriade 2000).

The study of this phonological process has the potential to provide some insight the precarious status of allophonic vowel nasalization in dialects of English which lies somewhere on the continuum from a phonetic or mechanical effect to fully phonemicized with phonologized in-between. Along this diachronic process of sound change, the cue to nasality shifts from the following consonant to the vowel which facilitates the eventual "destruction" of the phonetic precursor (Hyman 1976). In general American English, this process has obviously not fully taken place yet since nasality is not contrastive on vowels; however, vowels before nasal consonants have been demonstrated to be nasalized to a degree which must be phonological rather than purely phonetic in this language. Historically speaking, this is the typical path which leads to the emergence of phonemic nasal vowels in a given language and regressive spread is the more common mechanism (Ferguson 1975:181).

Therefore, I believe it is valuable to examine how variable coda nasal consonant deletion rules apply in such a process where phonologization is underway and which we know from diachrony, has the potential to result in phonemicization. I hypothesize that, compared to speakers of general American English, speakers of African American English will rely more readily on coarticulatory vowel nasalization as a cue to a following nasal consonant regardless of its presence or absence. This is because, for these speakers, the coarticulatory nasalization as a cue to the following nasal consonant licenses its deletion due the redundancy of auditory-acoustic cues. The experimental result should manifest as earlier discrimination on a continuous [la] ∼ [lã] ∼ [lãn] scale among participants who are speakers of African American English or have experienced significant exposure to AAVE.

In a broader sense, the experiments proposed in this paper is designed to ask: what is the underlying representation of coarticulatorily nasalized vowels and how do speaker groups with different linguistic backgrounds differ in their perceptual mapping from a continuous auditory-acoustic signal to abstract phonetic categories? The difficulty of the question lies in the constellation of cues which are at play—to be elaborated further in the background section, but which implicates lingual articulations—in perception of nasality and the primacy of auditory-acoustic cues in relation to one another in different contexts. When cues are in conflict, does categorical perception fail? This would suggest that when cues conflict it is the gesture that is invariant in perception and non-cooperating cues are uninterpretable. The listener must untangle whether an acoustic percept is a result of a lingual or velic gesture. To this end, vowel height will constitute an additional manipulation in the experimental design with the goal of relating results of the perceptual study to typological observations regarding the emergence of phonemic nasal vowels cross-linguistically.

2 Background

In order to approach coda nasal consonant deletion, I first examine the acoustic cues available in the speech signal to a listener in determining the presence of a nasal coda consonant and the related coarticulatory vowel nasalization. In particular, I explore a possible relationship between the perceptual saliency of vowel nasality as a function of vowel height. This approach is situated within an effort to understand how perception has influenced sound patterns diachronically and synchronically. In a study of sound deletion or reduction, it's important to consider the factors which license deletion; that is, why should a speaker map something from the input to nothing in the output (i.e., violate the faithfulness constraint Max in Optimality Theory)? Or in other words, what conditions influence when is a contrast retained? To this question, Coetzee (2004:222) says: "The idea is that a contrast is preserved (licensed) more easily in contexts where the cues for its perception are more salient than in contexts where these cues are less salient." Deletion of a coda nasal consonant may be licensed when the cues to its presence are difficult to perceptually distinguish from nasalization on the preceding vowel. The assumption is that listeners, particularly AAVE speakers, attribute vowel nasalization to following tautosyllabic nasal consonants to the point that the presence of an actual stop closure constitutes redundant information; in fact, place of articulation information is already degraded for nasal stops due to the lack of a stop burst.

2.1 Coarticulatory Vowel Nasalization

The crucial point which serves as the impetus for this study rests in the observation that vowels preceding nasal consonants in English are nasalized as a consequence of overlapping vocalic and velic gestures for the vowel and the nasal consonant respectively (i.e., the velum is lowered early in anticipation of the following nasal consonant). During the articulation of a nasal consonant, air continues to flow through the nasopharyngeal port which results in nasal and sinusoidal (also called "paranasal") cavity resonances. As a result of continued airflow, the nasal stop consonants /m, n/ do not exhibit a clear stop burst because nasal airflow prevents an increase in intraoral pressure (Thomas 2011).

In the case of coarticulatory vowel nasalization in American English, real-time magnetic resonance imaging suggests its degree extends beyond an automatic or mechanical explanation in coda (CVN) position, but not for onsets (NVC), in American English (Byrd et al. 2009). Experimental evidence from a motor planning study which manipulated speaking rate provides further confirmation that coarticulatory vowel nasalization is allophonic rather than mechanical in nature in American English (Solé 2007). That is, as duration of a prenasal vowel increased in their study, the proportion of nasalization on the vowel increased in duration whereas the oral portion did not. To contrast, the duration of the nasalized portion of a vowel before nasal consonants in Spanish was static relative to increasing overall vowel duration indicating a mechanical or purely phonetic effect. In this language, coarticulatory nasalization occurred only as a result of purely mechanical effects stemming from the timing of the sluggish velic gesture overlapped with a lingual articulation.

Having established that vowels are phonologically nasalized preceding nasal consonants, I now discuss the actual acoustic and perceptual consequences. Acoustically speaking, nasalized oral vowels are characterized by a distinct prominence at F1 (A1) which is replaced by a "broader prominence consisting of two peaks" (Chen 1996:40). Based on numbers proposed by Chen (1997:2362), the paranasal (P0) "pole-zero pair [introduces] a 3.1 dB increase at 270 Hz if the bandwidth of the pole and zero are 120 Hz. If the bandwidths are 80 Hz, the pole-zero pair would introduce a 5.5-dB increase at 290 Hz." Other studies place the first paranasal formant within a comparable range from 250 to 300 Hz (Feng and Castelli 1996, Maeda 1993).

Because the perceptual value of formants decreases monotonically with increasing frequency, I will keep things simple by considering only A1 and P0 as they are the most perceptual salient. This is further motivated by the observations of Hawkins and Stevens (1985:pg. 1560–1561) that individual differences in nasal and paranasal cavity sizes make higher frequency ($\geq$ 1500 Hz) spectral consequences of nasality highly variable. We will also observe in the next section that this is motivated by results of perceptual studies of vowel nasalization and nasality.

2.2 Perception of Vowel Nasality: Vowel Height

According to ideas regarding quantal theory within acoustic phonetics, non-linearities between articulation and acoustic output produce contrasts in speech because they allow some articulatory "squishiness" on either side of an s-shaped curve (Stevens 1989). That is, articulations can be continuous and variable, but by crossing a certain threshold, speech sounds make a perceptual "leap" into a new "category." This principle can help reveal in what phonological environments a speaker might be motivated to violate faithfulness to allow a vowel to categorically consume the cue to nasality from a following nasal consonant. The interaction between oral (A1) and paranasal (P0) resonant frequencies in the F1 region provides ripe territory for examining differential perceptual cues of this sort because the (para)nasal resonances remain in fixed positions but first oral formant (F1) covaries with vowel height. I will examine this interaction in greater detail shortly, but it's important to establish the reason for doing so in the first place. Namely, the robustness of a cue should determine whether or not the nasal consonant is more likely to be implicated in a phonological process of deletion (Wright 2004). A more robust cue to nasality on the preceding vowel is predicted to fortify the odds of deleting a following nasal.

Returning to the s-shaped relationship between articulation and acoustic output, acoustic coupling leading to convergence between the first oral and paranasal formants provides guides for phonetic classification, but these configurations are not intrinsically stable and require a "perceptual mechanism" to "freeze" variations in the convergence zone (Schwartz et al. 1997:258–259). This has been explained with the observation that two formants which are critically close are perceived as a single prominence; this is termed the center-of-gravity effect and the critical distance has been identified as about 3.5 Bark (Christovich and Lublinskaya 1979). The question, then, is what F1 and P0 (first paranasal formant) configurations are the most and least likely to lead to formant convergence via the center-of-gravity effect?

In order to determine for which vowels Bark-converted P0 and A1 differences ($A1 - P0$) are theoretically at or below the critical 3.5 Bark distance (i.e., the formants are perceptually converged or "frozen" via a center-of-gravity effect), I used F1 Bark values given by Schwartz et al. (1997:pg. 266) in their paper on Dispersion-Focalization Theory of vowel systems and a P0 value of 270 Hz (2.72 Bark) as provided by Chen (1997).[1] Results based on these theoretical values indicate that only the paranasal formant of low vowels /a, ɑ, ɒ, a̟/ would be perceptually distinguishable from F1 and all other vowels, F1 would be perceived as a weighted average between A1 and P0. The distance between the A1 and P0 of [æ] is very near this perceptual threshold (3.34 Bark), so I include it as a low vowel in my analyses. As distance between P0 and A1 decreases with height, the high vowels /i, y, ɨ, ɯ, u/ are very close in $A1 - P0$ values and have a theoretical difference of only 0.18 Bark or 7 Hz.

Compared to empirical findings, these theoretical values appear to hold up. For example, consider the following evidence from articulatory as well as perceptual researcher. In articulatory terms, the high vowels, [i] and [u] in English exhibit significantly greater velar height and velopharyngeal contact than the low vowels [æ] and [ɑ] and with [n] showing the greatest velopharyngeal opening of all consonants (Moll 1962). Low, more open or sonorous vowels have an intrinsically lowered velum due to the effects of the levator-paliti muscle which is attached to both the velum and the tongue body (Gick et al. 2013). In an articulatory study, individuals were shown to be more prone to make lingual rather than velic compensation for the high vowel [ĩ] than they were for the low vowel [ã] (Carignan et al. 2011). This behavior is explained as a result of the effects of vowel nasality (namely the introduction of P0) which can be reanalyzed by the listener as vowel height (i.e., F1) because the first oral and paranasal formants have perceptually converged. These results reveal that individuals are able to perceptually distinguish vowels whose nasalization they can attribute to coarticulatory effects of the following nasal consonant (i.e., non-low vowels) from those where nasality is less obviously a property of the following nasal consonant (low vowels). Finally, these results are compatible with a series of perceptual and machine learning experiments reported by Styler (2015) which identified $A1 - P0$ and its corresponding bandwidth reduction as a primary cue for vowel nasality among English and French speakers. Incidentally, it was also reported that human listeners responded more accurately and quickly to stimuli with nasal codas than they did to nasal onsets which is exactly what is predicted by the information contained within coarticulatory vowel nasalization in English Styler (2015:pg. 134).

Despite the extensive research on perception and articulation of nasal and nasalized vowels, the complexity of the auditory-acoustic cues means much is still not yet known.

[1] Choosing other values between 250 and 300 Hz did not drastically alter the results.

In particular, the relationship between vowel height (F1) and the perception of nasality is relatively underexplored:

> Studies have not shown whether vowel height affects perceived nasality independent of intrinsic velic height and degree of spectral prominence. Given that the main effect of tongue body height is to shift the frequency of F1, and that the main effect of vowel nasalization is also in the region of F1, it may be that shifts in F1 frequency alone influence perceived nasality. . . . [S]ystematic study of the perceived nasality of manipulating F1 frequency has not been undertaken" (Beddor 2014:pg. 178)

Therefore, in addition to answering questions about the effects of linguistic experience on perception of coarticulatory vowel nasalization, this study will attempt to contribute to understanding about vowel height and nasality.

3 Categorical Perception: Identification Tasks

Categorical perception is the experience of discontinuity over a continuously changing series of stimuli across a category boundary, accompanying an absence of a clearly perceived change within a category (McQueen 1996). Categorical perception is tested using an identification (labeling) task paradigm in which stimuli from a physical continuum span two categories which are unambiguous at their end-points, e.g., [b] versus [p] on a voicing continuum. Participants are presented with two stimuli from different steps in the continuum and asked to identify which end of the continuum a given token belongs to. This task results in a labeling function which establishes at what stimulus point a continuum shifts abruptly in labeling probability from one category to another in perception (i.e., location of a category boundary or 50% crossover point). Between the endpoints of the continuum, the stimuli should be ambiguous in some way with respect to their category membership. This will be discussed in the elaboration of the stimuli design in the methods section.

Past findings in categorical perception using labeling and discrimination tasks report that increasing the period between presentation of stimuli results in increased probability that stimuli will be identified as the same sound when presented with within-category stimuli. Additionally, the inclusion of a word-like context increases sensitivity to vowel categories, for example when discriminating [bil], [bɪl], and [bɛl] versus these vowels in isolation. For this reason, the present experiments will be designed to avoid as many experimental confounds as possible in order to focus directly on the question at hand. Not every speech sound exploits psycho-acoustic thresholds, but the hope is that because categorical perception of vowel nasality has been reported, it can be used as a cue to the presence or absence of a following nasal consonant coda to which it may be attributed. Indeed, the 50% cross-over point in the identification function should shift in correspondence with the particular cue or cues relied on by the two groups of participants (Repp and Libermant 1984).

3.1 Non-native Perceptual Contrasts

Because the aim of this paper is to investigate whether differences in linguistic experience have an effect on the nature of the identification function, it is relevant to examine some background literature on differences between native and non-native speech perception. Although the focus is on very similar dialects of the same language (indeed, sociolects), the manner in which phonemic contrasts are made (in production or perception) using auditory-acoustic information may differ across these groups. From a classic study on perceptual reorganization during language development and much subsequent work, it is widely accepted that ability to discriminate phonetic distinctions across speech sounds not in one's native repertoire declines early in life (Werker and Tees 1984). Additionally, although perceptual research is lacking in this subfield (though it is gaining interest), it is clear that auditory-acoustic cues for phonological differences between dialects are perceptually salient and allow listeners to categorize talkers of different regional dialects (Clopper and Pisoni 2004, Thomas 2002).

In terms of the dialect groups relevant to the present study, for speakers of AAVE, there exists a lexical contrast between the item [le] "lay" and [lẽ] ∼ [lẽn] "lane" whereas a lexical contrast only exists between [le] and [lẽn] for general American English speakers. Due to this difference in linguistic experience, it is expected that AAVE speakers will exhibit a more robust (earlier) change in probability of identification between [le] and [lẽ] on a continuum from [le] to [lẽn] even though nasal vowels are not phonemic in this language. Speakers of general American English, on the other hand, should be expected to place a category distinction between [lẽ] and [lẽn] rather than between [le] and [lẽ] on the same continuum. In this way, the *internal structure* of the phonetic categories is hypothesized to be different for the two groups of listeners (Miller 1997). With respect to the present study, for AAVE listeners, [lẽ] may be a satisfactory token of [len] "lane," whereas for a listeners who are speakers of general American English, the same token may be a better token of [le] "lay."

Furthermore, the research of Clopper (2004) demonstrates that linguistic experience in the form of exposure to certain dialect variants (or lack thereof) via regional mobility results in differences in how these dialect variants are perceptually categorized. This points to an additional advantage of a perceptual experiment for the current proposal over a production study; namely, it is hoped that mere exposure to coda nasal consonant deletion as an available variable rule in AAVE will sufficiently reshape African American listeners' perceptual categories. By conducting a perceptual experiment, it is possible to obviate the concerns which come along with a production study including the observer's paradox and limited access to speakers who are "high participators" in AAVE.

4 Method: Experimental Design

In this section, I sketch an outline of a perceptual experiment intended to test the hypotheses set forth above as supported by background literature on categorical perception and non-native contrasts. The experiment involves categorical perception of the presence or absence of a coda nasal consonant along three continua of vowel heights. A simple identification task will be used to model identification functions for each of the two speaker groups for the three 7-point vowel height continua. Stated explicitly, my hypotheses are:

H1: African American listeners as a group are predicted to identify a stimulus along the continuum as [lVn] earlier than general American English listeners will (e.g., a 50% cross-over point at 4 versus 5).

H2: There will be a main effect for vowel height such that [lɑn] will be selected earlier along its continuum than [lin] or [len] on their respective continua due to the robustness of coarticulatory cues.

4.1 Participants

Participants will be drawn from two main dialect groups: African American (Vernacular) English and general American English. Among the African American participants, it is preferred that they grew up in linguistic environment which afforded significant exposure to AAVE either within their a family or broader social network. Assuming this substantial exposure to the dialect features which characterize this sociolect, it is not absolutely necessary that the participants themselves engage in coda nasal consonant deletion. As for the speakers of general American English, no single regional or social dialect need be specified; however, significant experience with a dialect or language which has phonemic nasal vowels will be a criteria for exclusion from final analyses. As the study will be carried out in California, it is expected that speakers of California English will be the greatest represented dialect group among the participants. In order to determine that these conditions are controlled for, a language background questionnaire will be administered to all participants. The survey will include a question straightforwardly asking participants to identify themselves as speakers of Black English or not.

4.2 Stimuli

Stimuli will be synthesized from naturally produced utterances of [lV], [lṼ] and [lṼn] to create three continua from [l]-initial oral open syllables on one end to [l]-initial nasalized closed syllables ending in the alveolar nasal consonant [n]. Natural utterances in each of the frames for low, mid, and high vowels will be obtained from a male speaker of general American English because previous research on vowel nasality is largely based on knowledge of the male anatomy.

Each continuum will consist of seven levels for each of the three vowels: [ɑ], [e], and [i]. These configurations were chosen because lexical word status is identical for each end of each continuum. Finally, the first half of each oral-to nasal continuum will vary in degree of coarticulatory nasality on the vowel; the second half of each syllable will vary according to increasing duration of the alveolar nasal stop consonant [n] while maintaining the highest level of vowel nasalization. The continua are schematized as below using text size of the vowel and nasal consonant to represent increasing nasality and duration respectively:

"law"	[lɑ]	[lɑ̃]	[lɑ̃]	[lɑ̃]	[lɑ̃n]	[lɑ̃n]	[lɑ̃n]	"lawn"
"lay"	[le]	[lẽ]	[lẽ]	[lẽ]	[lẽn]	[lẽn]	[lẽn]	"lane"
"Lee"	[li]	[lĩ]	[lĩ]	[lĩ]	[lĩn]	[lĩn]	[lĩn]	"lean"

Table 1: Three stimulus continua for low, mid, and high vowels varying in degree of nasalization and duration of the coda nasal consonant.

Nasalized stimuli will synthesized by manipulating the amplitude of $A1-P0$ in naturally produced vowel tokens following the method reported by Styler (2015:pg. 122–123). This involves binning each vowel into three windows, measuring the naturally occurring $A1-P0$ in each window using a Praat script, and independently modifying the amplitude of F1 and the amplitude of $P0$. Using this technique, three levels of vowel nasalization will be produced: small, medium, and large magnitude. Greater degree of nasality corresponds to lower $A1-P0$ values. The duration of the alveolar nasal consonant will vary on the final three steps of the continua from 10ms to 50ms in 20ms steps. Small-sized steps are used in manipulation duration of the coda nasal consonant because this end of the continuum is of less interest to the purposes of the study.

In order to verify that the stimuli are reliably distinct within-category, the same script used to measure the natural nasality will be re-run on the synthesized stimuli and statistical analysis of $A1-P0$ values will ensure that nasalization levels are indeed statistically significantly different (i.e., is $A1-P0$ different between [lã] and [lɑ̃]?). Finally, it's possible that an AX discrimination task could be developed to further confirm that the within-category synthesized nasalization levels are perceptually distinguishable by human listeners. Following analyses of these results, the stimuli will either be adjusted accordingly or submitted to the identification task discussed in the next subsection.

4.3 Identification Task

The experiment will be implemented in PsychoPy software and will be carried out in the sound-attenuated booth in the UCLA phonetics lab. Stimuli from all three continua will be presented in nearly randomized order for each participant with ten repetitions of each stimulus token with one exception: the same stimulus will never be heard twice in a row. Each participant will therefore make identifications on 210 trials overall or 70 per vowel height continuum.

Study participants will be told that they will hear a single word once and be shown two written options (e.g., "lay" and "lane") on a computer monitor. Their task is to use designated buttons on a Cedrus button response pad to choose which word they heard (e.g., the red button will be labeled "left" and the blue button will be labeled "right"). In order to control for handedness and other factors, the appearance of the written words on either side of the screen will be randomized so that the oral and coda nasal ends of the continua will not always appear on the same side. Figure 1 below gives a representation of what a participant will see on a screen as they hear a stimulus, for example [lẽ] (stimulus #4):

Participants will be instructed to respond as quickly as possible on each trial as reaction times will be important for analyses of the responses. Participants will be given the option for a short break after the presentation of 70 stimuli in order to avoid fatigue. It is expected that the full identification task should not take more than 15 to 20 minutes.

5 Discussion of Analysis of Results

Following the collection of the data, analyses will be made on the dependent variable which is the relative probability of category identification at each stimulus number on the continuum.

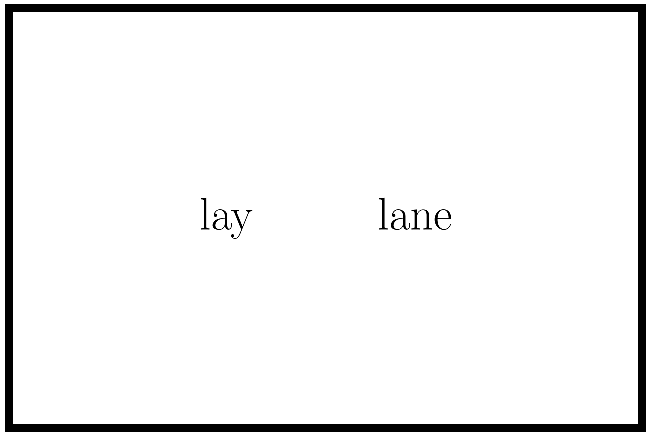

Figure 1: Example of a computer monitor screen during the experimental presentation of stimulus trials.

As a default, the probabilities will be calculated with respect to percentage identified as [lVn] (stimulus #7) as the reference level. In addition to stimulus number (1 through 7) as an independent variable, other independent variables include listener dialect group (AAVE or general American English) and the three levels of vowel height.

To test for the main effect of listener dialect background as well as influences of vowel height on probability of identifying a stimulus across the continuum as [lVn], a generalized linear mixed effects model will be constructed in `R` using the `glmer` function in the `lme4` package (Bates et al. 2014, R Core Team 2014). A logistic regression is appropriate for modeling these data in which the outcome variable is a binary choice as a function of independent predictor variables. Individual listener will be entered as a random effect in the model and likelihood ratio tests using the `anova` function will be employed to test for the significance of the independent variables.

5.1 Dialect Group Comparison

To illustrate a possible outcome of the identification experiment using a small example, consider a case in which H1 was indeed supported by the data. This would mean that listeners with AAVE dialect backgrounds assigned an earlier stimulus a significantly greater likelihood of [lVn] than did general American English listeners. This outcome is demonstrated in Figure 2 above with imagined data.

We can see from the identification functions in Figure 2 that the point at which stimuli are identified as [lãn] is earlier on average for listeners with AAVE linguistic background than it is for speakers of general American English. This result, if significant, would indicate that participants with significant exposure to AAVE more readily make use of coarticulatory vowel nasalization as a cue to a following nasal consonant, which is in turn what licenses its deletion in production. For general American English speakers, on the other hand, it may be less useful to consider vowel nasalization as a cue when it is consistently paired with a following nasal consonant in their language input. At the same time, it will certainly be more informative to this group of listeners than to listeners whose native language has no phono-

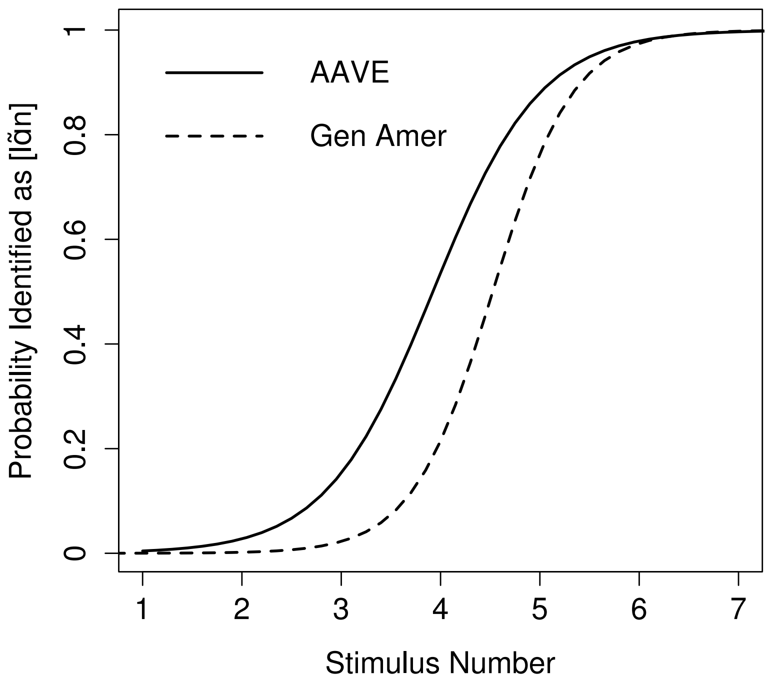

Figure 2: Illustration of identification functions for AAVE and general American English listeners.

logical vowel nasalization or phonemic nasal vowels (e.g., Spanish). A future experiment could explore this possibility by including more dialect/language experience groups.

Because the identification task may simply be too easy or not involve enough potential for within-category discrimination due to improperly synthesized stimuli, it will also likely be informative to closely examine reaction times obtained from each trial during the task. If the identification functions are completely overlapping between participants with AAVE backgrounds and general American English speaking participants, a result which shows that the former is able to make quicker judgments would still be compatible with H1.

If instead the results still fail to yield any difference, a final resort may be to modify the task to involve stimulus identification in noise. This would provide insight into confusability between stimuli on the continua which are otherwise readily recognized by both speaker groups when the experiment is made to be more tasking on the auditory-perceptual system. This manipulation would be undertaken by noise masking the coda nasal consonant which effectively alters the research question to: What degree of coarticulatory vowel nasalization is required for listener groups with different linguistic backgrounds in order to permit perceptual restoration of the absent phoneme in noise? Again, in this case phoneme restoration is expected to be possible with a lesser degree of vowel nasalization for participants with an AAVE background compared to general American English speaking participants. As before, it may be useful to additionally examine performance of Spanish-speaking monolinguals as a comparison group.

5.2 Effect of Vowel Height

Finally, by using three unique continua based on synthesized stimuli representing low, mid, and high vowels, it will be possible to test for connections between saliency of $A1 - P0$ as a function of F1 as predicted in H2. In this section I will explicitly outline the results which would align with predictions based on the articulatory and acoustic studies of vowel nasalization reviewed in the background. These results will be important for understanding vowel height and nasality cue relationships regardless of listener dialect (AAVE or general American English), though interesting interactions may also emerge.

First, I expect that 50% cross-over points in the identification functions (e.g., as represented in Figure 2) will be shifted towards earlier stimuli for the low vowel [ɑ] compared to non-low vowels [e,i]. This is because, as previously noted, $A1 - P0$ is perceptually distinguishable from F1 in low vowels, but not non-low vowels, and therefore attributable to the nasality introduced by coarticulation with a following nasal consonant. For non-low vowels, the introduction of a nasal pole-zero pair may be instead interpreted as a property of F1.

It is not clear in advance that this will necessarily be the outcome of the experiment in terms of vowel height effects which makes this an interesting proposal to carry out in the future. Many acoustic cues in addition to $A1 - P0$ play a role in perception of vowel nasality including vowel duration and formant bandwidths which are not being investigated here. For example, in an extension of this experiment, it would be valuable to additionally manipulate vowel duration as a variable which is associated with compensation for a deleted coda consonant, i.e., duration is longer for vowels preceding a deleted segment.

6 Conclusion

The results of this experiment would allow insight into the value of coarticulatory vowel nasalization as a cue to a following coda nasal consonant among groups with differing linguistic experience. Although these effects may turn out to be very subtle as both speaker groups have extensive experience with vowel nasalization associated with a following nasal consonant, it may yet be the case that actual coda consonant deletion has an impact on the relative value of this cue. Furthermore, I have outlined in the previous subsections various ways in which a small effect may be amplified using techniques such as phoneme restoration in noise.

This study could potentially have a more general impact on our understanding of the relationship between vowel height and coarticulatory vowel nasalization. Much room remains for the exploration of additional acoustic cues like vowel duration and formant bandwidths in future studies. Finally, one variable which was left unexplored in this proposal was the guise of the talker recruited to produce the stimuli to be synthesized. It should be determined what effect an African American guise versus a white American guise has on different dialect groups' identification functions if any. It may be that significant differences will only emerge when AAVE participants listen to an African American guise.

References

Bates, D., Maechler, M., Bolker, B., and Walker, S. (2014). *lme4: Linear mixed-effects models using Eigen and S4*. R package version 1.1-6.

Beddor, P. S. (2014). The perception of nasal vowels. In Anderson, S. R., Huffman, M. K., Krakow, R. A., and Keating, P. A., editors, *Nasals, Nasalization, and the Velum*, volume 5, pages 171–196. Elsevier.

Byrd, D., Tobin, S., Bresch, E., and Narayanan, S. (2009). Timing effects of syllable structure and stress on nasals: A real-time MRI examination. *Journal of Phonetics*, 37(1):97–110.

Carignan, C., Shosted, R., Shih, C., and Rong, P. (2011). Compensatory articulation in American English nasalized vowels. *Journal of Phonetics*, 39(4):668–682.

Chen, M. Y. (1996). *Acoustic correlates of nasality in speech*. PhD thesis, Massachusetts Institute of Technology.

Chen, M. Y. (1997). Acoustic correlates of English and French nasalized vowels. *The Journal of the Acoustical Society of America*, 102(4):2360–70.

Christovich, L. A. and Lublinskaya, V. V. (1979). The 'center of gravity' effect in vowel spectra and critical distance between the formants: Psychoacoustical study of the perception of vowel-like stimuli. *Hearing Research*, 1:185–195.

Clopper, C. G. (2004). *Linguistic Experience and the Perceptual Classification of Dialect Variation*. PhD thesis, Indiana University.

Clopper, C. G. and Pisoni, D. B. (2004). Some acoustic cues for the perceptual categorization of American English regional dialects. *Journal of Phonetics*, 32:111–140.

Coetzee, A. (2004). *What it means to be a loser: Non-optimal candidates in Optimality Theory*. PhD thesis, University of Massachusetts, Amherst.

Feng, G. and Castelli, E. (1996). Some acoustic features of nasal and nasalized vowels: A target for vowel nasalization. *The Journal of the Acoustical Society of America*, 99(6):3694–3706.

Ferguson, C. A. (1975). Universal Tendencies and 'Normal' Nasality. In Charles, F. A., Hyman, L. M., and Ohala, J. J., editors, *Nasalfest: Papers from a Symposium on Nasals and Nasalization*, pages 175–197, Stanford University: Language Universals Project.

Gick, B., Wilson, I., and Derrick, D. (2013). *Articulatory Phonetics*. Wiley-Blackwell, West Sussex, UK.

Hawkins, S. and Stevens, K. N. (1985). Acoustic and perceptual correlates of the non-nasal-nasal distinction for vowels. *Journal of the Acoustical Society of America*, 77(4):1560–1575.

Hyman, L. (1976). Phonologization. In Juilland, A., editor, *Linguistic Studies Presented to Joseph H. Greenberg*, pages 407–418. Anma Libri, Saratoga.

Labov, W. (1972). *Sociolinguistic Patterns*. University of Pennsylvania, Philadelphia.

Maeda, S. (1993). Acoustics of vowel nasalization and articulatory shifts in French nasal vowels. *Phonetics and Phonology*, 5:147–167.

McQueen, J. (1996). Phonetic categorisation. *Language and Cognitive Processes*, 11:655–664.

Miller, J. L. (1997). Internal structure of phonetic categories. *Language and Cognitive Processes*, 12:865–869.

Moll, K. L. (1962). Velopharyngeal closure on vowels. *Journal on Speech and Hearing Research*, 5:30–37.

R Core Team (2014). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Repp, B. H. and Libermant, A. M. (1984). Phonetic category boundaries are flexible. In Harnad, S. N., editor, *Categorical Perception*, pages 31–53. Cambridge University Press, New York.

Risdal, M. L. (2015). Phonology of nasal consonant deletion in African American English. Presented at UCLA Phonofest.

Schwartz, J.-L., Boë, L.-J., Vallée, N., and Abry, C. (1997). The dispersion-focalization theory of vowel systems. *Journal of Phonetics*, 25:255–286.

Solé, M. J. (2007). Controlled and mechanical properties of speech. In Solé, Beddor, and Ohala, editors, *Experimental Approaches to Phonology*, pages 302–321.

Steriade, D. (2000). Directional asymmetries in place assimilation: A perceptual account. In Hume, E. and Johnson, K., editors, *Perception in Phonology*. Academic Press.

Stevens, K. S. (1989). On the quantal nature of speech. *Journal of Phonetics*, 17:3–45.

Styler, W. (2015). *On the acoustical and perceptual features of vowel nasality*. PhD thesis, University of Colorado.

Thomas, E. R. (2002). Sociophonetic Applications of Speech Perception Experiments. *American Speech*, 77(2):115–147.

Thomas, E. R. (2011). *Sociophonetics: An introduction*. Palgrave Macmillan, London, UK.

Werker, J. F. and Tees, R. C. (1984). Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development*, 7:49–63.

Wolfram, W. (1989). Structural variability in phonological development: Final nasals in vernacular Black English. In Fasold, R. W. and Schiffrin, D., editors, *Language Change and Variation*, pages 301–322. John Benjamins, Philadelphia, Pennsylvania/Amsterdam, The Netherlands.

Wright, R. (2004). A review of perceptual cues and cue robustness. In Hayes, B., Kirchner, R., and Steriade, D., editors, *Phonetically-Based Phonology*. Cambridge.