Why So Short? Competing Explanations for Variation

Susanne Gahl
University of California at Berkeley

1. Introduction

"Please remain on the line. Your call is very ^{m}perdnus." As odd as this transliteration of "important to us" may look in print, the utterance rendered here struck its addressee as perfectly natural and intelligible in context. The actual pronunciation of words in continuous speech varies widely and tends to differ quite markedly from the citation forms of words. A large part of that variation is systematic, a fact that has made pronunciation variation, and phonetic reduction especially, the focus of much attention in Psycholinguistics, Phonetics, and Phonology.

Phonetic reduction is commonly taken to mean shortening of words and segments, deletion of segments and entire syllables, and articulatory undershoot resulting in lenition of consonants and centralization of vowels. Efforts to understand which forms are likely to undergo reduction, and under what circumstances, have yielded a number observations that are fairly uncontroversial: There is wide-spread agreement that high-frequency words are often phonetically reduced. Similarly, words that are "predictable", or especially likely to occur in a given context, are likely to undergo reduction. The reduction of the words "very important to us", given the context "Please remain on the line. Your call is ___" would be an example of reduction on the basis of predictability. Observations along these lines are well established and widely cited (Bell et al., 2003, Bell et al., 2009, Fowler et al., 1997, Hunnicutt, 1985, 1963). There is broad consensus that phonetic reduction is extremely common, that it represents a unitary phenomenon with many manifestations (durational shortening, lenition, deletion, centralization), and that this phenomenon calls for a unified explanation.

The consensus ends when attention turns to the reasons why reduction takes place: Do words undergo reduction because they are predictable from the point of view of the listener, or from the point of view of the speaker? Do speakers strengthen and lengthen certain words for the benefit of their listeners, and allow themselves to reduce words that they deem likely to be intelligible even when reduced? Or do speakers lengthen or shorten words depending on how much time they themselves need for planning and encoding their utterances?

The goal of the present paper is to summarize some recent work that asks what pronunciation variation in conversational speech reflects, by considering what happens when intelligibility and ease of production lead to different expectations about whether words are likely to undergo reduction or its opposite, phonetic strengthening. The lexical property at the focus of this work is phonological neighborhood density, a measure of how many other words in the lexicon are similar to a given target word. As explained in greater detail below, words in "dense" phonological neighborhoods tend to be challenging targets for recognition, other things being equal, but easy targets for production. If pronunciation variation reflects speakers' attempts to maximize intelligibility, words in dense neighborhoods should come to be strengthened and lengthened, to offset the detrimental effects of neighborhood density on recognition probability. If pronunciation variation reflects the speed and ease of production planning and encoding, then words in dense neighborhoods should come to be shortened and reduced. To preview the overall conclusions of that work: In conversational speech, regression models of word duration and vowel centralization suggest that words in dense neighborhoods are phonetically reduced, other things being equal.

---

[*] Some of this work is a synopsis of work described in detail in Gahl, Yao & Johnson (in press). Earlier versions of the analyses and detailed descriptions of the methods can be found in Yao (2011).

2. Explanations of phonetic reduction

Approaches to pronunciation variation have followed two basic strategies, each of them based on a different central observation. “Listener-based” approaches are based on the observation that predictable forms tend to be highly intelligible, even when reduced. This central observation motivates a general hypothesis about the relationship between intelligibility and articulation explanation consistent with the Hypo-and Hyperarticulation (“H & H”) theory articulated in Lindblom (1990). Reduction of predictable forms is consistent with that theory because, as stated in Flemming (2010), “speakers try to facilitate communication by speaking more clearly where the listener is likely to have greater difficulty with word recognition. Consequently, any factor that makes a word more difficult to recognize is expected to make a speaker more inclined to produce it clearly.” Conversely, on this view, speakers would be less inclined to produce words clearly that are easy targets for recognition due to their frequency or to their contextual predictability. Several proposals along these lines, and evidence consistent with this general line of explanation have been reported in the literature (e.g. Kilanski, 2009, Munson and Solomon, 2004, Munson, 2007, Scarborough, 2010).

The second basic strategy in explaining pronunciation variation has been to relate variation to speaker-internal mechanisms underlying language production, such as ease and speed of retrieval. The basic observation underlying such “production-based” approaches is that words that tend to get reduced are words that are easily retrieved. Ease and speed of retrieval of items from memory have been modeled in spreading-activation models (Anderson, 1983, Dell, 1986). In such models, likelihood of accurate retrieval, and lexical retrieval speed, are modeled as a function of “activation” of nodes in connectionist networks. High lexical frequency, for example, may be modeled as high resting activation of relevant nodes that enable high frequency items to be retrieved quickly and with great accuracy. Several proposals relating pronunciation variation to ease and speed of retrieval, and evidence consistent with this line of explanation, are reported in the literature (e.g. Gahl, 2008, Pluymaekers et al., 2005).

Both of these approaches may also be stated at a computational, as opposed to algorithmic, level of analysis, by relating patterns of variation to more global properties of utterances. Flemming (2010), for example, outlines an analysis relating a listener-oriented model of pronunciation to a Bayesian approach to word recognition. The central idea behind the proposal is that the degree of phonetic reduction is a function of the prior probability of a given speech signal being recognized correctly, i.e. categorized as an instance of the intended word (cf Jurafsky, 1996). Similarly, a number of proposals have sought to ground models of reduction in information-theoretic terms. The idea here is that reduction can be modeled as a function of information load: Predictable items, by definition, are low in information load and accordingly allotted little time and acoustic signal in utterances (Aylett and Turk, 2004, Jaeger, 2010, Levy and Jaeger, 2007).

Each of these approaches face certain empirical and conceptual challenges, and there is continued debate about their respective merits. The main empirical challenge for intelligibility-based accounts lies in the fact that the experimental record of settings in which speakers’ and listeners’ needs diverge provides strong evidence that there are severe limits on speakers’ ability to cater to their listeners’ needs (Bard and Aylett, 2005, Ferreira and Dell, 2000, Keysar, 2007). For example, one study (Bard & Aylett, 2005) found phonetic reduction in words that speakers had used before (in a map task), even when their addressees had changed. If phonetic reduction of repeated items in discourse were entirely due to speakers reducing items that are predictable from the point of view of their listeners, repeated words that are new to a given listener should not reduce, despite being repetitions as far as the speaker is concerned.

The main challenge to production-based approaches to phonetic variation concerns the relationship between ease of lexical retrieval and articulatory realization. The work reported in Gahl (2008), for example, is built on the assumption that variation in speaking rate in conversational speech reflects ease and speed of lexical access and retrieval: words that are easy to retrieve are reduced, because they are retrieved quickly. This explanation is attractive in that high-frequency words, for example, are retrieved quickly and tend to reduce. Empirically, this connection between fast retrieval and phonetic reduction is unassailable. However, conceptually the explanation is weak, since it seems to presuppose that speakers are always aiming to talk as fast as possible: After all, fast lexical leaves extra time that could be used to realize particularly ambitious articulatory goals: Why don't speakers realize articulatory targets accurately when they could do so without decreasing their overall speaking rate?

In fact, the assumption that high lexical activation, and therefore fast lexical retrieval, should be associated with phonetic reduction, is neither necessarily correct nor universally shared. A positive association between ease of retrieval/high lexical activation and the realization of extreme articulatory targets just outlined has in fact been assumed in explanations of certain types of phonetic variation. As one study asserts: "[H]igher activation levels […] lead to more active phonetic representations and consequently more extreme articulatory realizations." (Baese-Berk and Goldrick, 2009) High lexical activation, then, has been used to explain both phonetic reduction and phonetic strengthening.

Although each of these challenges is serious, they are probably not the main reason why there seems to be no end in sight to the debate between listener-based and production-based approaches to variation. The main reasons for the continued debate, in my view, are that (1) both are correct explanations of some relevant phenomena; and (2) these competing approaches make identical predictions in many empirical domains: High-frequency words, for example, tend to be easy targets for recognition, as well as for production. To move beyond the dichotomy, then, one needs to examine not so much which explanation is correct, but how these two sources of variation combine. Since many phenomena are attributable to speakers or listeners equally well at a common sense level, one also needs to examine critically those empirical domains in which the two types of explanations do make divergent predictions.

The goal of the work described here is to probe the respective roles of intelligibility and ease of production in phonetic reduction in conversational speech as a function of phonological neighborhood density, i.e. a lexical property that affects intelligibility and ease of production in ways yielding different predictions.

3. Previous literature

The main reason why listener-based and production-based explanations of variation make such similar predictions is that "for the most part, the same things that make a word easy to understand make that word easy to say" (Dell and Gordon, 2003). As pointed out in Dell & Gordon (2003), there is an interesting exception to that pattern: Phonological neighborhood density. Phonological neighborhood density is an index of the number of words that are phonologically similar to a target word. By one common criterion, words are considered phonological neighbors if they differ by insertion, deletion, or substitution of one phoneme. For example, the word *cat* has three lexical neighbors in the toy lexicon depicted in Figure 1 (*kit*, *cap,* and *map*), whereas the word *mop* only has a single neighbor in that lexicon (*map*).

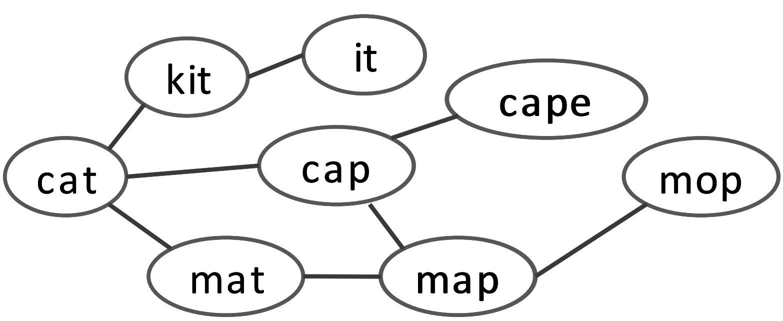

Figure 1. A toy lexicon

A key observation in models of word recognition is that recognition of words in "dense" phonological neighborhoods is poor, compared to recognition of words in sparse neighborhoods. Phonological neighbors appear to act as competitors (Luce and B., 1987, Luce et al., 1990, Luce and Pisoni, 1998). A more recent observation is that phonological neighbors act as supporters in production: There is a tendency for production of words in dense neighborhoods to be faster and more accurate, other things being equal (Gordon, 2002, Vitevitch, 1997, Vitevitch and Sommers, 2003).

These facts about phonological neighborhood density may seem puzzling at first, but are rooted in the nature of the tasks of listening and talking. Listeners encounter signals and must work out what meaningful units those signals represent. The main thing that can go wrong for listeners is that they mistakenly categorize a signal as representing a different meaningful unit – and this is most likely to happen when there are words that sound similar to an intended target. Unsurprisingly, then, phonological competitors are the main sources of word-level errors in recognition. Speakers, on the other hand, begin their task with some idea of what meanings they wish to convey and subsequently select suitable forms. The main source of word-level errors (as opposed to segmental errors) in production are semantic competitors (Dell et al., 1997). A speaker might intend to say "feline" and say "canine" by mistake. The phonological competitor "beeline" represents a far less likely error in production, less likely than a listener mishearing "feline" for "beeline". (Of course, a speaker might end up making a segmental error resulting in "beeline", after all.) In the two-step interactive spreading-activation models of lexical retrieval described in Dell & Gordon (2003), once a semantic target is selected, it receives additional activation through its phonological neighbors: As a result, not only do phonological neighbors not threaten production targets the way they do recognition targets, they in fact act as supporters.

Given these divergent effects of phonological neighborhood density on recognition and production, it is clear that listener-based and production-based explanations of pronunciation variation make different predictions about the effect of this variable on phonetic reduction. A listener-based account of pronunciation variation would lead one to expect high phonological neighborhood density to be associated with phonetic strengthening: Speakers should, on this view, strengthen words in dense neighborhoods, to improve their intelligibility. A production-based account, on the other hand, would lead one to expect high neighborhood density to be associated with phonetic reduction: Just like high word frequency, high phonological neighborhood density makes words easy to retrieve, other things being equal, and efficient retrieval is associated with phonetic reduction.

What are the effects of phonological neighborhood density on pronunciation? One acoustic parameter that has been examined in this connection is vowel dispersion (or its opposite, vowel centralization): Vowel dispersion represents the distance between a vowel token and the center of vowel space. Figure 2 shows a schematic representation of the place of three cardinal vowels in vowel space. The center of that space is marked with a schwa:

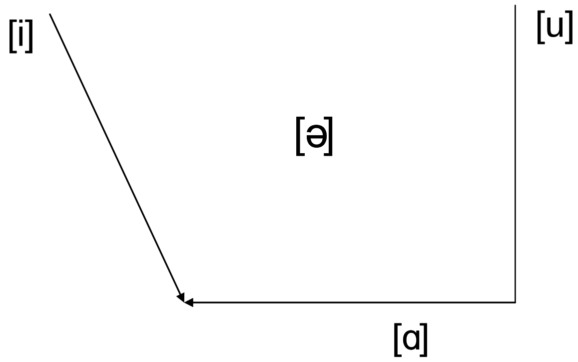

Figure 2. A toy vowel space

Of course, in reality, vowel tokens are far more widely scattered through this space. Figure 3 shows a more realistic vowel space, plotting tokens of CVC words produced by one talker in the course of a single interview by the first two vowel formants. The interview is part of the Buckeye corpus of conversational speech (Pitt et al., 2005, Pitt et al., 2007).

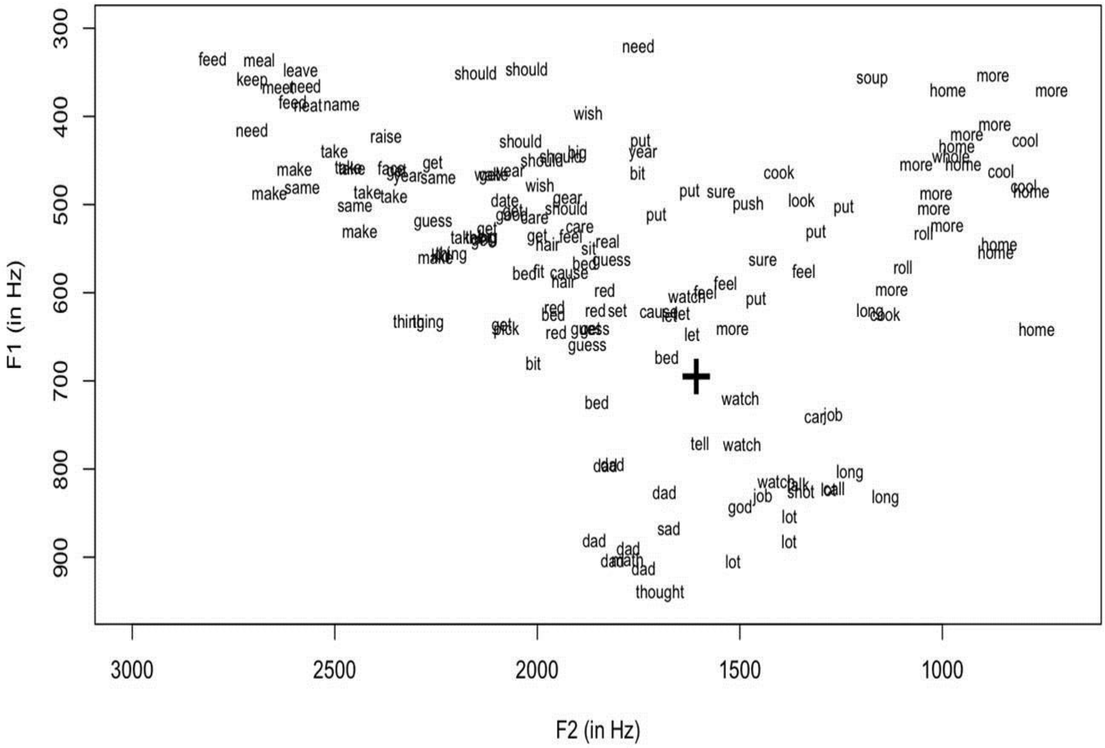

Figure 3. An actual vowel space

As can be seen in Figure 3, tokens representing a given phoneme are widely dispersed throughout vowel space, and the areas associated with different phonemes overlap. The center of the talker's vowel space, estimated as the average F1/F2 of tokens of schwa, is marked with a plus sign.

Vowel dispersion affects intelligibility (Bradlow et al., 1996) and is among the variables talkers manipulate when asked to maximize intelligibility (Moon and Lindblom, 1994, Picheny and Durlach, 1985), making it a prime candidate for an aspect of pronunciation variation that might reveal the role of intelligibility in pronunciation variation.

Several studies have examined vowel dispersion as a function of neighborhood density. The first study to do so was an investigation of this variable (Wright, 1997, Wright, 2004) in a database of recordings of ten talkers (Torretta, 1995). Wright (1997, 2004) found that vowels were significantly more centralized in high-frequency, low-density words than in low-frequency, high-density words. Since word frequency and neighborhood density were jointly manipulated in the database, the results do not indicate which of these variables was responsible for the observed effect. Several subsequent studies report similar patterns, however, of high neighborhood density being associated with phonetic strengthening, rather than its opposite. A subsequent study (Munson and Solomon, 2004) probed the effects of word frequency and neighborhood density by factorially manipulating these two variables, also in a single-word naming task: It was found that low frequency and high density were each associated with increased vowel dispersion relative to high frequency and low density. Similar patterns of increased dispersion, and other acoustic properties associated with increased intelligibility, for words in dense neighborhoods are also reported in a number of other studies (e.g. Baese-Berk and Goldrick, 2009, Kilanski, 2009, Peramunage et al., 2010, Scarborough, 2005, Scarborough, 2010, Watson and Munson, 2007).

Two things are striking about these studies. The first is that they do not replicate the usual association of high word frequency and durational shortening (not all of the studies mentioned looked for this association, of course). The second is that, without exception, they are based on experimental tasks requiring speakers to produce word lists or words in short carrier phrases (such as "Say __ to me again" or "The first word is __, the word after ___ is ___".). These two facts may be related: When producing word lists, talkers tend to read words at a regular pace. This tendency has been found to override effects of lexical properties, such as word frequency, on timing. Several of these studies do reproduce other, non-temporal effects of word frequency, finding high word frequency to be associated with decreased vowel dispersion and other manifestations of articulatory undershoot.

The temporal characteristics of word lists and elicited phrases, the effects of stimulus ordering may be interacting with effects of neighborhood density. To investigate the effects of this variable in a setting that might shed light on the kinds of phonetic reduction under discussion in the literature mentioned above, we therefore investigated the effects of this variable in conversational speech.

4. Methods

We examined the effect of phonological neighborhood density on two aspects of phonetic realization: word duration and vowel dispersion, using mixed-effects regression models with Word and Talker as random effects and a range of known determinants of vowel dispersion and word duration. An early version of that study can be found in Yao (Yao, 2011). The theoretical implications of the study, and a description of the control variables and modeling procedure, can be found in Gahl (Gahl et al., 2011, in press).

All data came from the Buckeye Corpus of conversational speech (Pitt, et al., 2007; Pitt, Johnson, Hume, Kiesling, & Raymond, 2005), which consists of ca. one hour of spontaneous speech from each of 40 talkers (20 male, 20 female) from Columbus, Ohio.

We analyzed vowel dispersion and word duration in CVC monomorphemic content words in the corpus. The corpus contained 534 word types that met our inclusion criteria for the analyses of word durations, represented by 12,414 tokens from 40 talkers. The exclusion of central vowels and diphthongs meant that the set of words in the analysis of vowel dispersion was a subset of the words in the analysis of word durations, consisting of 414 word types, represented by 9,075 tokens from 39 talkers. The two sets of words were analyzed in two separate models.

5. Results

The relationship of the control variables to word duration and vowel dispersion was what one would expect, given previous studies: For example, increasing frequency, bigram probability given neighboring words, and speaking rates were associated with shorter word durations. The proportion of variability accounted for was low compared to some previous models of word and segment duration in connected speech (Bell, et al., 2009; Gahl, 2008; Quené, 2008). This difference is likely to be due in part to the fact that the studies just cited included utterance-final and pre-pausal tokens. Phrase-final position and disfluencies produce large effects on word duration, making it possible to account for a substantial portion of variability in duration based on these two predictors alone. Figure 4 shows the partial effects of all fixed effects in the final models of word durations and vowel dispersion. Crucially for the point of the study, increased neighborhood density was associated with shorter word durations and reduced vowel dispersion.

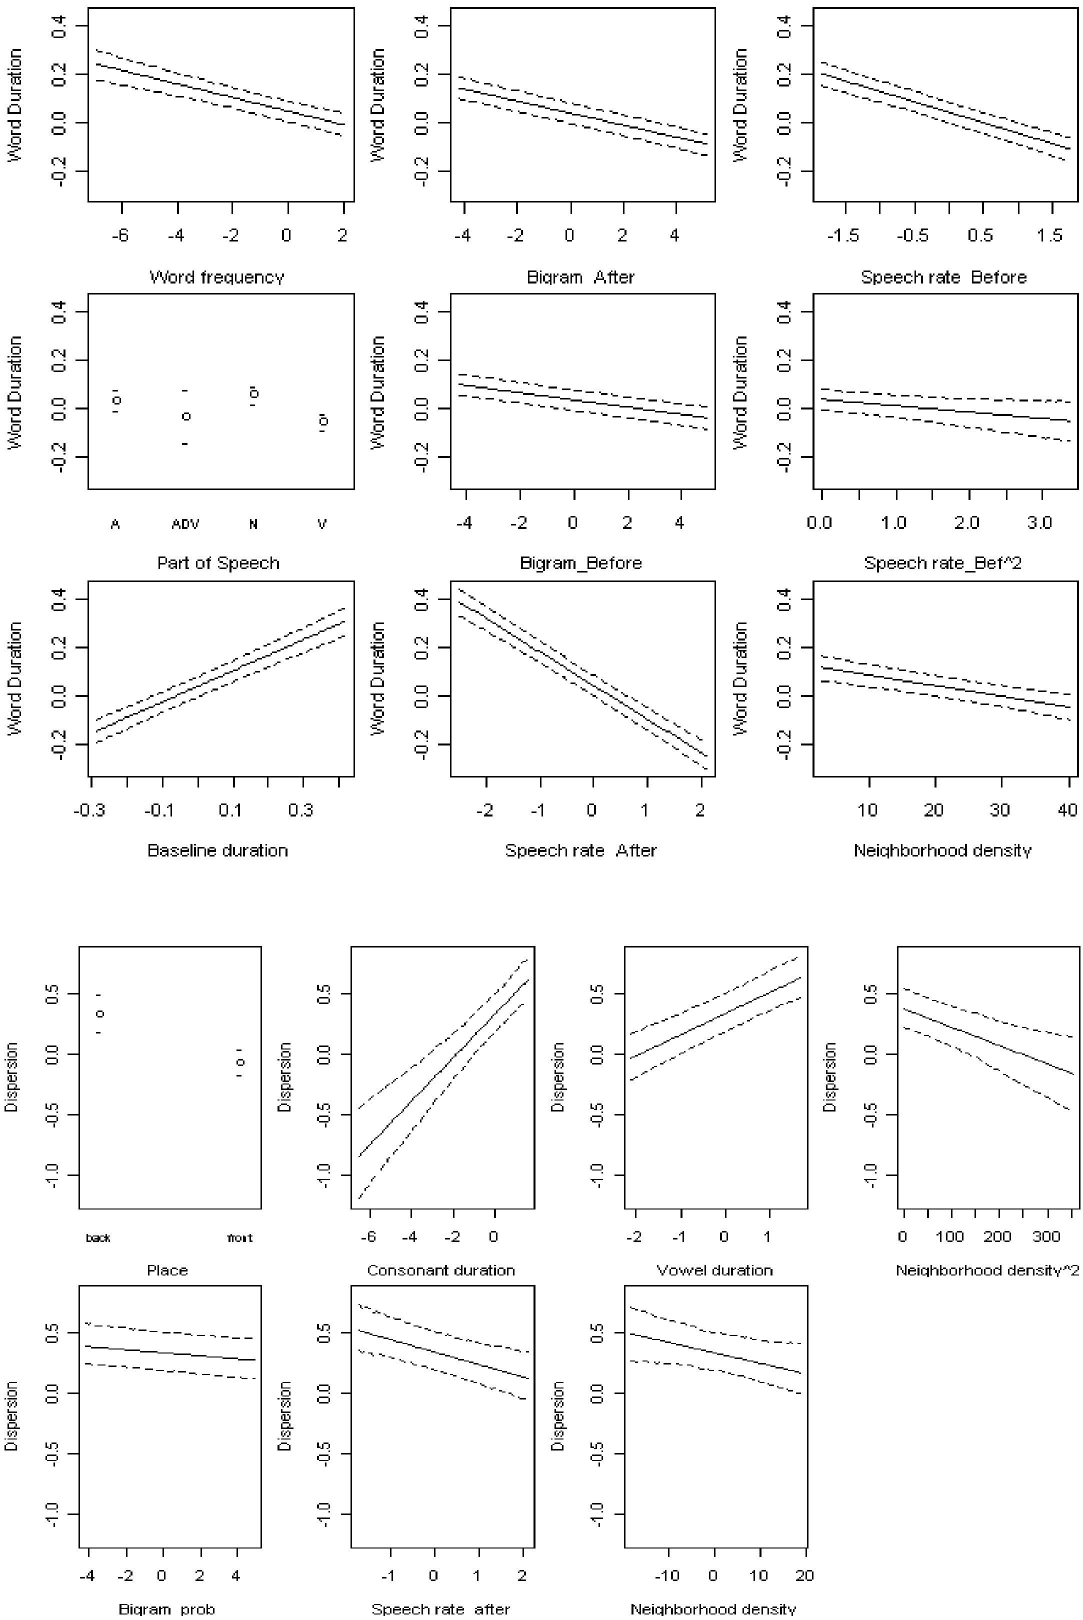

Figure 4. Partial effects in a model of word duration and vowel dispersion

6. Discussion

The central finding in the study just described is that high phonological neighborhood density was associated with phonetic reduction (shorter word durations and decreased vowel dispersion) in a corpus of conversational speech. This finding is unexpected under accounts of pronunciation variation that attribute variation to speakers' attempts to maximize intelligibility of words for the benefit of their listeners. The finding is also unexpected given previous results showing high phonological neighborhood density to be associated with acoustic properties that increase intelligibility, such as increased vowel dispersion.

I have hinted at one possible reason for the apparent discrepancy between the pattern observed in the corpus and the earlier, experimental findings: When people are asked to read lists of words, they tend to read at a regular pace. When words are produced at a regular pace, fast lexical retrieval leaves more time for speakers to realize extreme articulatory targets. As has been argued elsewhere (Pierrehumbert, 2001), words in dense phonological neighborhoods may have more extreme articulatory targets than words in sparse neighborhoods. High-frequency words, on the other hand, may have less extreme (more 'reduced') articulatory targets. This difference in articulatory targets could yield the observed pattern: When temporal and attentional constraints allow it, e.g. when speakers are producing words at an even pace and are facing few other attentional demands, pronunciations reflect articulatory targets.

The temporal demands of conversational speech are quite different from list reading. In particular, speakers reading words one at a time don't need to be planning any upcoming utterances, but are concerned only with the current trial. Perhaps the working assumption that fluent conversational speech reflects speakers' speaking as fast as possible isn't far off the mark, after all: Articulation rate certainly lags behind the speed with which speakers can think of things to say, and speakers may for all intents and purposes be trying to realize words as quickly as they become available for articulation.

7. Conclusion

Listener-based and production-based explanations are mutually compatible: It is almost certainly the case that multiple sources of variation are responsible for the range of linguistic variation. There are clear instances of variation attributable to one source of variation or the other: Disfluencies in the face of word-finding difficulties, or slurred speech of intoxicated talkers might be extreme examples of variation due to production-based factors; whereas foreigner-directed speech, baby talk, or elder speak uncontroversially involve speakers' adapting their speech to what they perceive to be appropriate in talking to a particular audience. A large and continually growing literature in Sociophonetics attests to talker adaptation to particular social and linguistic settings and particular addressees. An equally vast literature on the limits of Audience Design, and on characteristics of speech produced under conditions that impede language production attests to the fact that variation in speech reflects many processes not under talkers' control and not conducive to successful communication.

These clear types of variation leave much else to be accounted for. The reduction of "very important to us" in the example above needn't be caused by the same mechanisms that give rise to different speaking styles or the mechanisms that cause us to slow down in the face of word-finding difficulties. Yet, in the face of broad consensus that "predictable" items tend to be phonetically reduced, it is important to ask what makes items "predictable": predictability from the point of view of the listener, or the speaker? The goal of the work summarized here was to probe the respective roles of intelligibility and ease of production in phonetic reduction in conversational speech, in a class of cases in which intelligibility and ease of production make divergent predictions. More broadly, the goal is to address the relationship between single-word production and conversational speech. Single-word psycholinguistic experiments have yielded a tremendous amount of insight into the lexical properties that affect language production. Research on language production in conversational speech can profit from that insight, but only given an increased understanding of the differences between the timing of processes underlying single-word production and connected speech.

References

Anderson, John R. (1983). Spread of Activation. Journal of Experimental Psychology: Learning, Memory, and Cognition 10, 791-798.

Aylett, Matthew, and Turk, Alice (2004). The Smooth Signal Redundancy Hypothesis: A Functional Explanation for Relationships between Redundancy, Prosodic Prominence and Duration in Spontaneous Speech. Language and Speech. Language and Speech 47, 31-56.

Baese-Berk, Melissa, and Goldrick, Matthew (2009). Mechanisms of interaction in speech production. Language and Cognitive Processes 24, 527-554.

Bard, Ellen Gurman, and Aylett, Matthew P (2005). Referential form, word duration, and modeling the listener in spoken dialogue. In Approaches to studying world-situated language use, eds. John C. Trueswell and Michael K. Tanenhaus, 173-191. Cambridge, MA: MIT Press.

Bell, Alan, Jurafsky, Daniel, Fosler-Lussier, Eric, Girand, Cynthia, Gregory, Michelle, and Gildea, Daniel (2003). Effects of disfluencies, predictability, and utterance position on word form variation in English conversation. Journal of the Acoustical Society of America 113, 1001-1024.

Bell, Alan, Brenier, Jason, Gregory, Michelle, Girand, Cynthia, and Jurafsky, Dan (2009). Predictability effects on durations of content and function words in conversational English. Journal of Memory and Language 60, 92-111.

Bradlow, Ann R., Torretta, Gina, and Pisoni, David B (1996). Intelligibility of normal speech I: Global and fine-grained acoustic-phonetic talker characteristics. Speech Communication 20, 255-272.

Dell, Gary S (1986). A spreading-activation theory of retrieval in sentence production. Psychological Review 93, 283-321.

Dell, Gary S., Schwartz, Myrna F., Martin, Nadine, Saffran, Eleanor M., and Gagnon, Deborah A (1997). Lexical access in aphasic and nonaphasic speakers. Psychological review 104, 801-838.

Dell, Gary S., and Gordon, Jean K (2003). Neighbors in the lexicon: Friends or foes? In Phonetics and Phonology in Language Comprehension and Production, eds. Niels O. Schiller and Antje S. Meyer, 9-47. New York: Mouton De Gruyter.

Ferreira, Victor S., and Dell, Gary S (2000). Effect of ambiguity and lexical availability on syntactic and lexical production. Cognitive Psychology 40, 296-340.

Flemming, Edward (2010). Modeling listeners: Comments on Pluymaekers et al. and Scarborough. In Laboratory Phonology 10, eds. Cécile Fougeron, Barbara Kühnert, Mariapaola D'Imperio and Nathalie Vallée, 587-606. Berlin: Mouton De Gruyter.

Fowler, Carol A., Levy, Elena T., and Brown, Julie M (1997). Reductions of spoken words in certain discourse contexts. Journal of Memory & Language 37, 24-40.

Gahl, Susanne (2008). "Time" and "thyme" are not homophones: Word durations in spontaneous speech. Language 84, 474-496.

Gahl, Susanne, Yao, Yao, and Johnson, Keith (2011, in press). Why reduce? Phonological neighborhood density and phonetic reduction in spontaneous speech. Journal of Memory and Language in press.

Goodman, Nelson (1943). On the simplicity of ideas. Journal of Symbolic Logic 8, 107–21.

Gordon, Jean K (2002). Phonological neighborhood effects in aphasic speech errors: Spontaneous and structured contexts. Brain and Language 82, 113-145.

Hunnicutt, Sharon (1985). Intelligibility vs. redundancy - conditions of dependency. Language and Speech 28, 47-56.

Jaeger, Florian (2010). Redundancy and reduction: Speakers manage syntactic information density. Cognitive Psychology 61, 23-62.

Jurafsky, Daniel (1996). A probabilistic model of lexical and syntactic acces and disambiguation. Cognitive Science 20, 137-194.

Keysar, Boaz (2007). Communication and miscommunication: The role of egocentric processes. Intercultural Pragmatics 4, 71-84.

Kilanski, Kelley (2009). The effects of token frequency and phonological neighborhood density on native and non-native speech production, Linguistics, University of Washington.

Levy, Roger, and Jaeger, Florian (2007). Speakers optimize information density through syntactic reduction. In Twentieth Annual Conference on Neural Information Processing Systems.

Lieberman, Philip (1963). Some effects of semantic and grammatical context on the production and perception of speech. Language and Speech 6, 172-187.

Lindblom, Björn (1990). Explaining phonetic variation: A sketch of the H&H theory. In Speech production and speech modeling, eds. William J. Hardcastle and Alain Marchal, 403-439. Dordrecht: Kluwer.

Luce, Paul A., and B (1987). Neigborhoods of words in the mental lexicon: Dissertation/Thesis.

Luce, Paul A., Pisoni, David B., and Goldinger, Steven D (1990). Similarity neighborhoods of spoken words. In Cognitive models of speech processing: Psycholinguistic and computational perspectives., ed. Gerry T.M. Altmann, 122-147. Cambridge, MA, US: The MIT Press.

Luce, Paul A., and Pisoni, David B (1998). Recognizing Spoken Words: The Neighborhood Activation Model. Ear & Hearing 19, 1-36.

Moon, Seung-Jae, and Lindblom, Björn (1994). Interaction between duration, context, and speaking style in English stressed vowels. Journal of the Acoustical Society of America 96, 40-55.

Munson, Benjamin, and Solomon, Nancy Pearl (2004). The Effect of Phonological Neighborhood Density on Vowel Articulation. Speech, Language, and Hearing Research 47, 1048-1058.

Munson, Benjamin (2007). Lexical Access, Lexical Representation, and Vowel Production. In Laboratory Phonology 9, Phonology and Phonetics, eds. Jennifer Cole and José Ignacio Hualde, 201-227. Berlin: Mouton de Gruyter.

Peramunage, Dasun, Blumstein, Sheila E., Myers, Emily B., Goldrick, Matthew, and Baese-Berk, Melissa (2010). Phonological neighborhood effects in spoken word production: An fMRI study. Journal of cognitive neuroscience 23, 593-603.

Picheny, Michael A., and Durlach, Nathaniel I (1985). Speaking Clearly for the Hard of Hearing I. Journal of Speech, Language, and Hearing Research 28, 96-103.

Pierrehumbert, Janet B (2001). Exemplar dynamics: Word frequency, lenition and contrast. In Bybee, Joan (Ed); Hopper, Paul. (2001). Frequency and the emergence of linguistic structure. Typological studies in language, vol. 45., 137-157. Amsterdam, Netherlands: John Benjamins Publishing Company.

Pitt, Mark A., Johnson, Keith, Hume, Elizabeth, Kiesling, Scott, and Raymond, William (2005). The Buckeye corpus of conversational speech: Labeling conventions and a test of transcriber reliability. Speech Communication 45, 89-95.

Pitt, Mark A., Dilley, Laura, Johnson, Keith, Kiesling, Scott, Raymond, William, Hume, Elizabeth, and Fosler-Lussier, Eric (2007). Buckeye Corpus of Conversational Speech (2nd release). Columbus, OH: Department of Psychology, Ohio State University (Distributor).

Pluymaekers, Mark, Ernestus, Mirjam, and Baayen, Harald (2005). Lexical frequency and acoustic reduction in spoken Dutch. Journal of the Acoustical Society of America 118, 2561-2569.

Scarborough, Rebecca A (2005). Coarticulation and the Structure of the Lexicon (Doctoral dissertation, UCLA Department of Linguistics, 2004). Dissertation Abstracts International A 66, 574.

Scarborough, Rebecca A (2010). Lexical and contextual predictability: Confluent effects on the production of vowels. In Laboratory Phonology 10, eds. Louis Goldstein, Douglas H. Whalen and Catherine T. Best, 557-586. Berlin, New York: De Gruyter Mouton.

Shannon, Claude E. (1948). A mathematical theory of communication. Bell System Technical Journal 27, 379–423, 623–656.

Shosted, Ryan (2006). Correlating complexity: A typological approach. Linguistic Typology 10, 1–40.

Torretta, Gina (1995). The "easy-hard" word multi-talker speech database: An initial report. In Research on Spoken Language Processing Progress Report, 321-333. Bloomington, IN: Speech Research Laboratory, Indiana University.

Vitevitch, Michael S (1997). The neighborhood characteristics of malapropisms. Language and Speech 40, 211-228.

Vitevitch, Michael S., and Sommers, Mitchell S (2003). The facilitative influence of phonological similarity and neighborhood frequency in speech production in younger and older adults. Memory & Cognition 31, 491-504.

Watson, Peter J., and Munson, Benjamin (2007). A comparison of vowel acoustics between older and younger adults. Proceedings of the 16th International Congress of Phonetic Sciences, 561-564.

Wright, Richard (1997). Lexical competition and reduction in speech: A preliminary report. Indiana University Research on Spoken Language Processing Progress Report No. 21, 471-485.

Wright, Richard (2004). Factors of lexical competition in vowel articulation. In Papers in Laboratory Phonology VI, eds. J. Local, R. Ogden and R. Temple, 26-50. Cambridge: Cambridge University Press.

Yao, Yao (2011). Effects of Neighborhood Density on Pronunciation Variation, Linguistics, University of California at Berkeley.

Proceedings of the 29th West Coast Conference on Formal Linguistics

edited by Jaehoon Choi, E. Alan Hogue, Jeffrey Punske, Deniz Tat, Jessamyn Schertz, and Alex Trueman

Cascadilla Proceedings Project   Somerville, MA   2012

Copyright information

ISBN 978-1-57473-451-5 library binding

Ordering information

Orders for the library binding edition are handled by Cascadilla Press.
To place an order, go to www.lingref.com or contact:

Cascadilla Press, P.O. Box 440355, Somerville, MA 02144, USA
phone: 1-617-776-2370, fax: 1-617-776-2271, sales@cascadilla.com

Web access and citation information

This entire proceedings can also be viewed on the web at www.lingref.com. Each paper has a unique document # which can be added to citations to facilitate access. The document # should not replace the full citation.

This paper can be cited as:

Gahl, Susanne. 2012. Why So Short? Competing Explanations for Variation. In *Proceedings of the 29th West Coast Conference on Formal Linguistics*, ed. Jaehoon Choi et al., 1-10. Somerville, MA: Cascadilla Proceedings Project. www.lingref.com, document #2681.