What is LabPhon? And where is it going?

Janet B. Pierrehumbert, Northwestern University
Cynthia G. Clopper, Ohio State University

Abstract

A look back at the first ten Laboratory Phonology conferences suggests three primary contributions of our field to the study of language. First, an examination of the citation patterns of a sample of laboratory phonologists revealed that our work has influenced research in such disparate areas as linguistics, psychology, computer science, and medicine. Second, the original LabPhon interest in the relationship between mind and matter has led to the advancement of autosegmental-metrical phonology and provided for a growing body of research on the relationship between abstract linguistic structures and their phonetic implementation. Third, the nature of linguistic representations themselves has been challenged and forced us to think about our assumptions regarding segmental and prosodic structures. As we look forward to the next twenty years, we suggest that the future of laboratory phonology lies in understanding another component of the embodied linguistic system: the social function and social context of language.

As we mark the end of the first dectuple of Laboratory Phonology conferences and look forward to the second dectuple, Cynthia and I are honored to reflect with you on what laboratory phonology has become and where it is going. In the United States, graduations are typically marked by "commencement" ceremonies as students begin the next chapter of their lives. Thus, this is our commencement address as the laboratory phonology community "graduates" to its second dectuple of meetings.

1. Who is LabPhon?

We were inspired by the question Abby posed in her chapter "What is LabPhon?", her retrospective look at the evolution of Laboratory Phonology since its inception, and her empirical approach to figuring out where the field began and what it has become. We would like to begin our discussion of the 10th anniversary session by considering a second, related, question, "Who is LabPhon?", by tabulating the scientific fields represented by publications produced by laboratory phonologists and the citations of those publications across disciplines. It is said that you will be known by your works and by the company you keep, and for scientists, these two properties come together in publication and citation patterns. We examined the citation record of a representative set of laboratory phonologists to see just how our works are known and what kind of company we are keeping. We approached this question in an unusual way with the assistance of Stefan Wuchty, a postdoctoral fellow of the Northwestern Institute on Complex Systems. Stefan suggested creating a citation network to visualize the scientific fields that laboratory phonologists have contributed to both directly (through their own publications) and indirectly (through the citation of those publications by other scientists).

We selected a set of 23 heavily cited laboratory phonology authors from the approximately 250 scholars who have published in the first nine laboratory phonology volumes and who presented papers at the 10^{th} meeting in Paris in 2006. The tabulation was conducted using the ISI Web of Science, an online searchable database of publications and citations for 5000 journals in a broad range of scientific disciplines. The focus of our investigation was the ISI field codes assigned to each publication included in the index. For readers who may be unfamiliar with the ISI Web of Science, (1) illustrates the type of record we were working with, for an article by Ken Stevens that exemplifies the type of research that made laboratory phonology possible. The Stevens and House (1955) study included 32 references and has itself been cited 159 times.

(1) Title: Development of a quantitative description of vowel articulation
Author(s): Stevens KN, House AS
Source: Journal of the Acoustical Society of America 27 (3): 484-493 1955
Document Type: Article
Language: English
Cited References: 32 Times Cited: 159
Subject Category: Acoustics

For our set of 23 laboratory phonologists, Stefan tabulated the ISI subject category field codes indicating broad research areas for each paper and each of its citations. The resulting estimates of publication and citation records for these 23 authors are conservative given that the ISI Web of Science does not include edited volumes or conference proceedings in its database. However, together, the 23 selected authors have produced 367 papers in ISI Web of Science journals under 34 different ISI subject category field codes. The 367 papers have, in turn, been cited by 8928 papers under 115 field codes. Figure 1 shows the number of citations in each of the 20 most frequent field codes for this set of papers. Not surprisingly, the top four fields are applied linguistics (which includes journals such as *Language and Speech*), experimental psychology, acoustics, and language and linguistic theory. However, the other fields include a range of research areas in the social sciences, medicine, and education, and reflect the total impact of laboratory phonology across a wide range of disciplines.

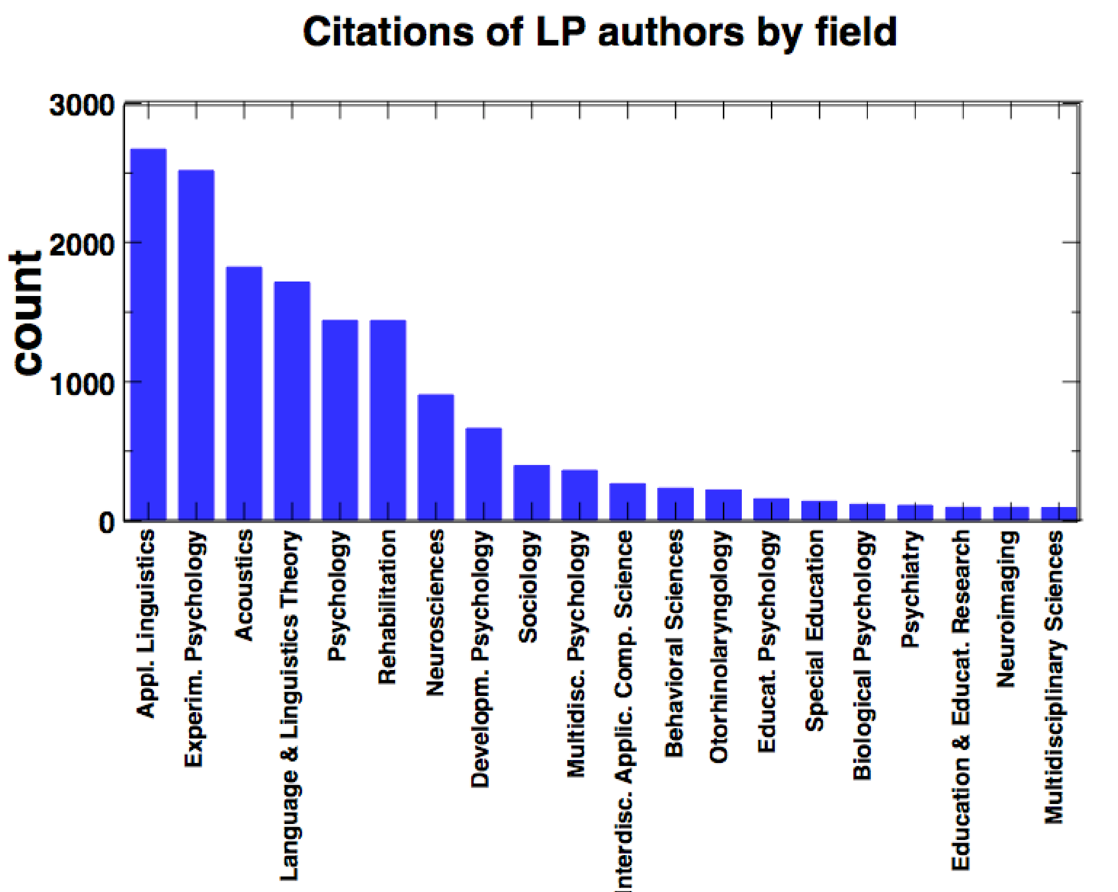

Figure 1. Number of citations in each of the 20 most frequent field codes for the 367 papers produced in ISI Web of Science journals by 23 representative laboratory phonologists.

The interdisciplinarity of laboratory phonology is even more evident when we construct a citation network linking papers published in one field to the fields in which those papers were cited. The nodes in the network are the field codes. For each of the 8928 citations, a link was created between the field code of the cited paper to the field code of the citing paper. In Figure 2, we have plotted the strongest 5% of these connections, with each link representing at least 135 cross-field citations. The thickness of the lines reflects the strength of the connection, with stronger links appearing as darker lines than weaker links. For the core areas (acoustics, applied linguistics, experimental psychology), most of the links are bidirectional. That is, cross-citation occurs frequently for these fields. However, some of the more peripheral fields show only unidirectional connections, suggesting the transport of knowledge across disciplines.

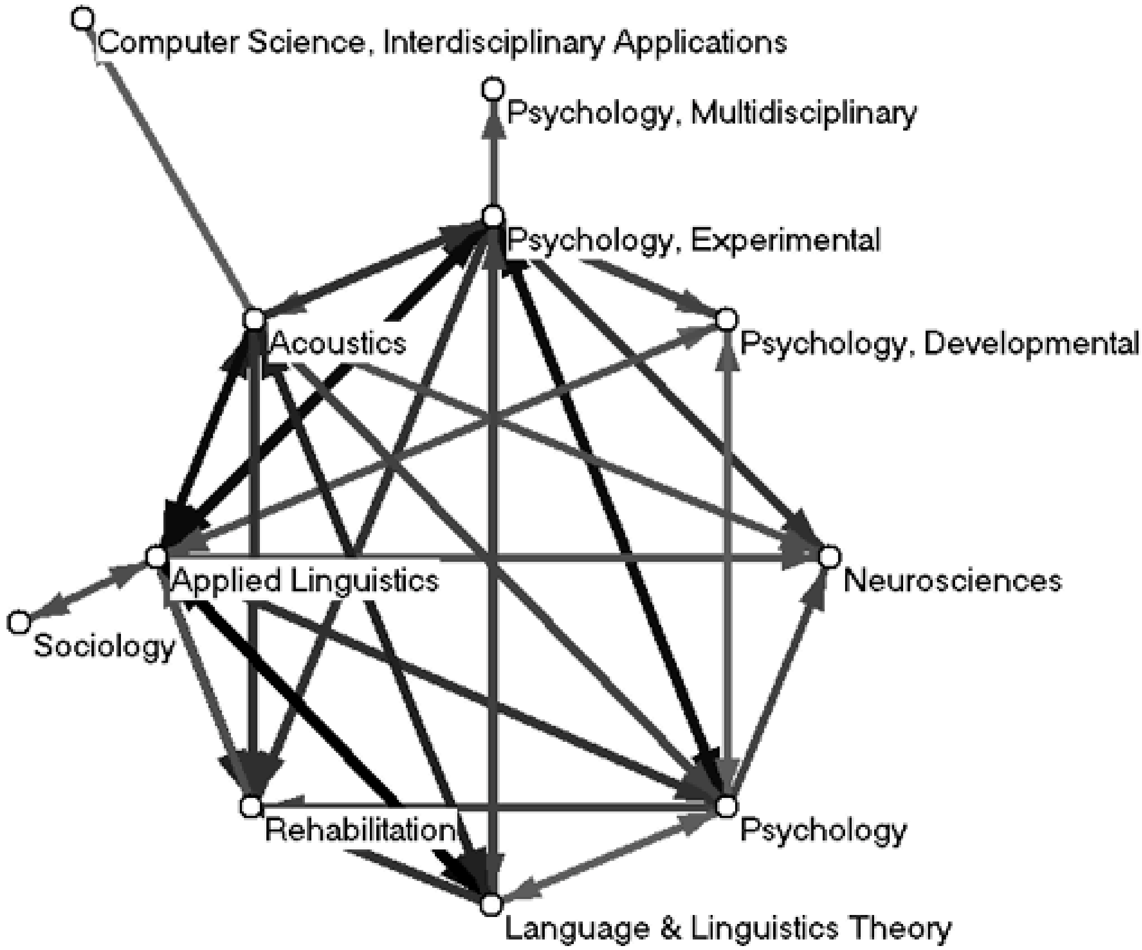

Figure 2. Laboratory phonology citation network: strongest 5% of the connections.

In Figure 3, we have plotted the strongest 20% of these connections, with each link representing at least 14 cross-field connections. We again observe an orbit around acoustics, applied linguistics, language and linguistic theory, and experimental psychology that reflects the core focus of laboratory phonologists. However, as Figure 3 demonstrates, laboratory phonologists have also made contributions to biomedical engineering, electrical and electronic engineering, and computer science. In addition, research in laboratory phonology has influenced such varied fields as radiology, zoology, pediatrics, developmental biology, physiology, and sport sciences, even though as a community we don't publish directly in those fields.

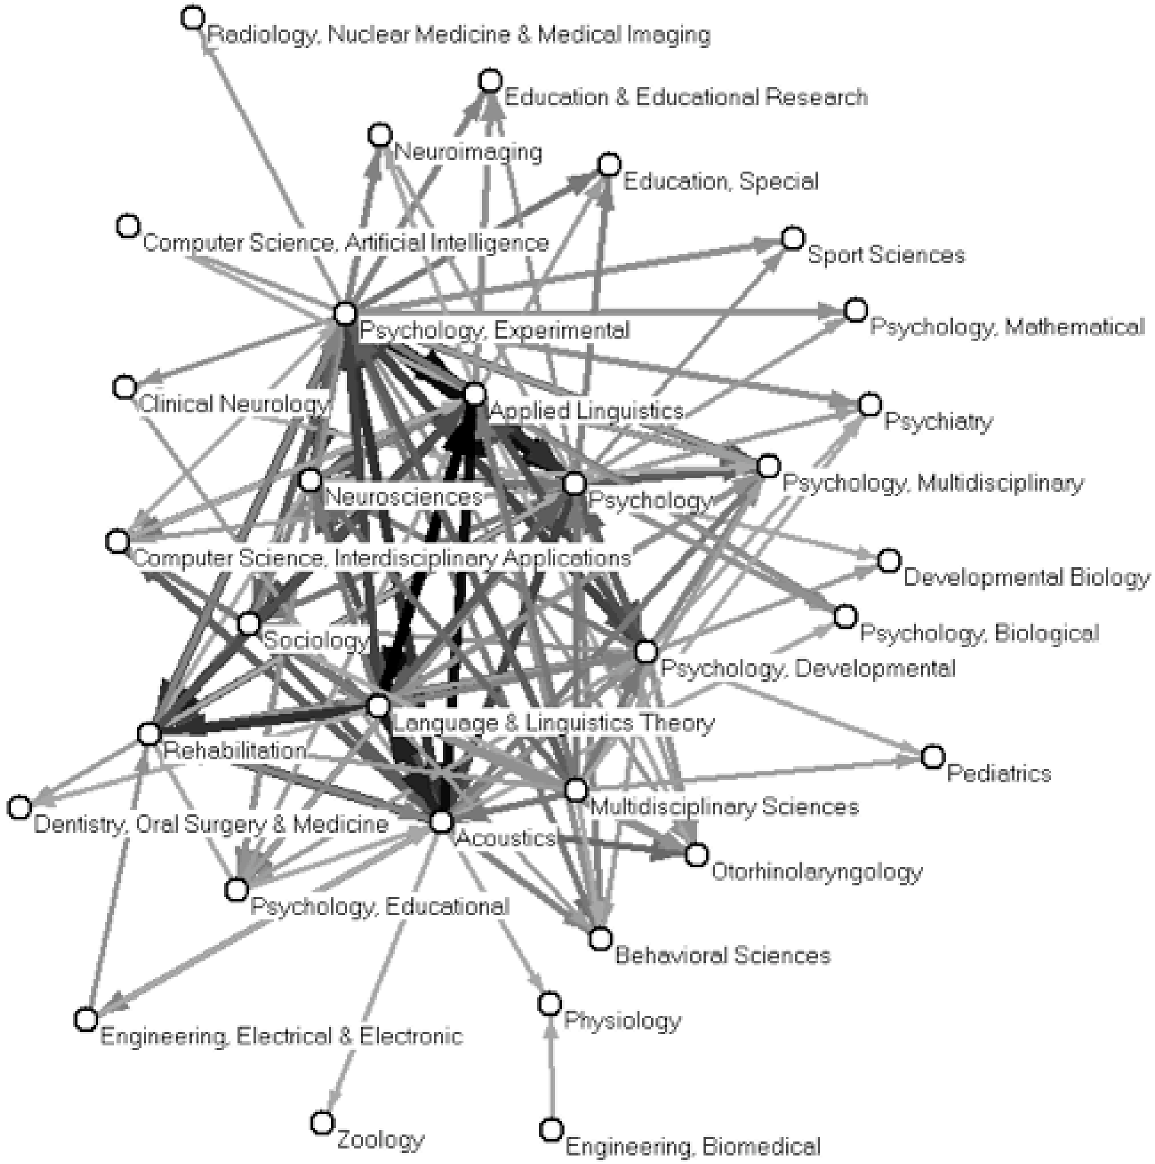

Figure 3. Laboratory phonology citation network: strongest 20% of the connections.

Thus, laboratory phonology is not a subfield of linguistics or cognitive science, but a federation of scholars with similar research interests and goals. Laboratory phonology is a free-trade zone and like all free-trade zones, it has important commodities that hold these long distance connections together. The original commodity that made laboratory phonology possible is the attunement of mind to matter; that is, the relationship between the physical reality of speech and the cognitive representation of language. In philosophy, this relationship is known as embodiment; in psychology, it is discussed in terms of the relationships among representation, perception, and action. This commodity is exemplified by the work of early speech scientists who were interested in the relationships between articulatory and acoustic phenomena.

From the very first Laboratory Phonology meeting, a second commodity has emerged that ties the community together: autosegmental-metrical phonology. Two of the papers in the 10th anniversary session (Byrd and Choi, Ridouane) drew from the autosegmental-metrical phonology line of thought. We will discuss each of these papers in turn, highlighting their contributions to the commodities of laboratory phonology and asking the reader to consider some further questions that these studies raise. We propose that these papers also suggest a new commodity for laboratory phonology: the relationship between levels of representation and the social function and social context of language. The last paper in the 10th anniversary session (Cutler et al.) dealt explicitly with levels of representation and we will use our discussion of that paper as a launching point for a discussion of the embodiment of language in a social system.

2. Mind and matter

The foundation for laboratory phonology was laid by the previous generation of researchers who recognized the importance of the link between mind and matter in language and speech. A landmark for thinking about this mind-to-matter question with respect to speech and language is Denes and Pinson's (1963) high school textbook, *The Speech Chain*. In the very first figure in their book, Denes and Pinson laid out the perception-production loop that characterizes the transmission of speech "from the brain of the speaker to the brain of the listener" (Denes and Pinson 1993: 5). In the speech chain model, sound is a communication channel among individuals and perception and production are on equal footing. Production transforms ideas into physical actions and perception transforms physical events into ideas. Their model also critically includes two ears, which allows the speaker to also serve as a listener for the purposes of self-monitoring, and two brains, which allows the talker and the listener to switch roles. Thus the speech chain crucially involves not only a perception-production loop, but also the communicative function of language in a social context.

Denes and Pinson did not refer explicitly to embodiment or the mind-to-matter problem in their text, but they did lay the groundwork for thinking about the relationships among articulation, acoustics, auditory processing, and abstract linguistic representation. The identification of speech as an embodied system meant that early speech scientists conducted foundational research that serves as a benchmark for what a scientific theory of embodied cognition should look like. For example, Ken Stevens' research on the relationship between articulation and acoustics and Georg von Bekesy's Nobel Prize in Physiology and Medicine in 1961 for his work on the anatomy of the inner ear provide a baseline against which laboratory phonology must be evaluated. We have the luxury of being brief and concluding our introductory remarks on embodiment here thanks to Jean-Luc Schwartz's extensive and insightful review of the mind-to-matter problem and the biological affordances of the production-perception loop in his commentary on an earlier session from the Paris meeting.

3. Autosegmental-metrical phonology

From the very first LabPhon conference, laboratory phonologists, both using and contributing to the formalism of autosegmental-metrical phonology, took the general concept of embodiment into a new domain. The original work on the mind-to-matter problem in speech and language dealt with phoneme inventories, whereas dynamical properties played no significant role. In the autosegmental-metrical approach, however, hierarchical data structures from different time scales are analyzed in a single theoretical framework. The physical reality of time and the abstract representations of segmental and prosodic structures are brought together under a single roof. Specifically, autosegmental-metrical phonology provided a unified representational ontology for segmental and tonal features. Consider, for example, Figure 4, adapted from Broe (1992), which shows a schematic of the parallels between contour and spreading phenomena for tones and segments. With respect to contours, we observe parallel structures for contour tones, affricates, and pre-nasalized stops. The spreading of tones is also seen as a parallel structure to long vowels and geminate consonants.

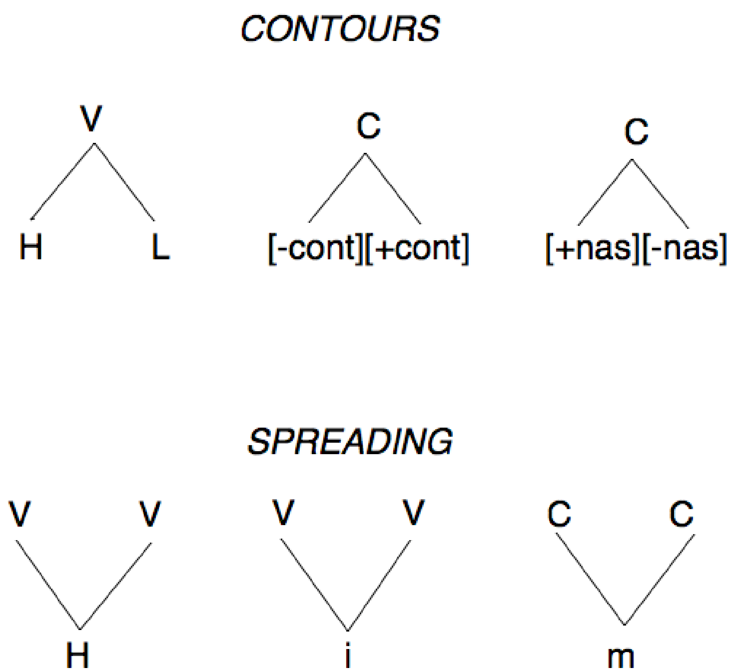

Figure 4. Contour and spreading phenomena in tones (left) and segments (center and right). Adapted from Broe (1992).

Autosegmental-metrical phonology also provides a unified ontology for lexical and phrasal structure. For example, Pierrehumbert and Beckman (1988) proposed that phonetic realization uses a running window on an autosegmental-metrical structure as shown in Figure 5. In this figure, the arrow indicates the time window of interest and the question is: in this time window, how is the phonological vowel-vowel hiatus executed phonetically? Pierrehumbert (1994) argued that prosodic nodes have canonical realizations in their own right. That is, they don't just induce hyper- or hypo-articulation of segmental contrasts, but instead prefer certain phonetic implementations of segments. In the example in Figure 5, the strong word *octane* “wants” a consonantal attack (i.e., glottal stop insertion). Then we can ask: how strong of a glottal stop is it? The strength of the glottal stop can depend on the structural description of the target, all the way up to the Intonation Phrase node. Here we assume that the window can “see” just the vowel-vowel hiatus and all structures above them, although the precise size of the window can vary. Aspiration and glottalization can similarly be accounted for by considering phrasal and lexical structure in parallel with phonetic implementation (see Fougeron and Keating 1997; Pierrehumbert and Talkin 1992).

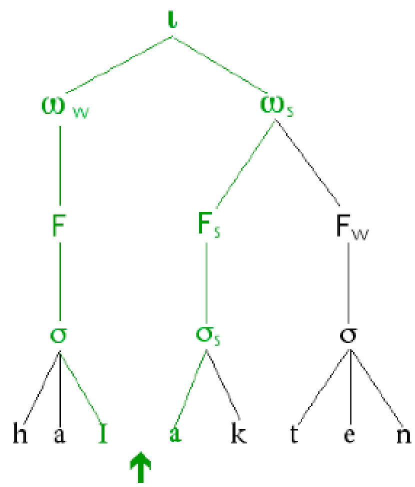

Figure 5. Autosegmental-metrical structure of *high octane*.

Two of the papers presented in the 10th anniversary session in Paris exemplify the underlying themes of autosegmental-metrical phonology and were undoubtedly selected by the organizing committee for the special 10th anniversary session as a result of their connection to this foundational component of laboratory phonology. We begin our discussion of these papers with Byrd and Choi, who presented a classic laboratory phonology production study examining the interaction between higher-level (phrasal boundary type) and lower-level (syllable position) metrical structure. This study follows from early laboratory phonology research on the relationship between prosodic constituency and the phonetic implementation of segments at prosodic boundaries (e.g., Beckman and Edwards 1990; Beckman, Edwards, and Fletcher 1992), as well as earlier work on the timing of sequences of articulatory gestures (e.g., Saltzman, Löfqvist, and Mitra 2000).

One of the primary contributions of Byrd and Choi's work is that it explored the impact of prosodic structure on articulatory relations, not just articulatory parameter values. They observed gradient differences in the timing of consonant gestures across different boundary types, which reflect the relations between consecutive gestures and not categorical timing parameters. Thus, abstract prosodic structures have a graded, non-categorical impact on physical, articulatory gestures.

A second important contribution of Byrd and Choi's study is that it explored individual talker variation. They found evidence of both intra-talker and inter-talker variation. As shown in their Figure 9, individual talkers exhibited the effects of phrasal boundary type on consonant gesture overlap to different extents for different cluster types. In addition, the three talkers differed from one another on the same dimensions. The question that remains unanswered, however, is what this variation tells us about the relationship between phrasal boundary type and consonant gesture overlap. That is, what does it mean for something to be statistically robust at the population level in the face of intra- and inter-talker variability across different cluster types? The role of individual differences within and across populations is an important component of linguistic systems, both in acquisition and for adults navigating a range of social contexts. The empirical nature of laboratory phonology research puts us in an excellent position to begin tackling these difficult questions.

Ridouane's analysis of the phonetic correlates of gemination continues a rich history in laboratory phonology of research on the structure and phonetic implementation of geminate consonants, including the phonetics and phonology of gemination resulting from assimilation (Ohala 1990), the articulatory gestures differentiating singleton and geminate consonants (Smith 1995), and the acoustic correlates of lexical and derived geminates (Ladd and Scobbie 2003). The case of Berber is a welcome addition to this literature because, most unusually, geminates are contrastive initially and finally as well as medially. Ridouane's study is notable for its use of electropalatography to make up for the limitations of acoustic measures, and well as for its integrated discussion of functional and formal aspects of the problem.

Ridouane's results clearly support his assertion that geminates arising across a word boundary involve two distinct structural positions and two distinct tokens of the melodic element. This conclusion follows from the spirantization facts and the lack of any effect on the preceding vowel. Thus, we may consider these geminates as pseudo-geminates. Though some other languages would release the first consonant in the sequence, we infer from Ridouane's discussion that Berber does not. The pseudo-geminates contrast with true geminates, in which a melodic element is spread across two structural positions. These occur contrastively in monomorphemic lexical items, or within complex words across morpheme boundaries. Their occurrence within complex words is understandable both formally (as a reflex of the Obligatory Contour Principle; see discussion in McCarthy 1986) and psycholinguistically (as a reflex of the fact that complex forms are listed in the lexicon if sufficiently frequent; see discussion in Hay 2003).

The comparison of singleton stops to pseudo-geminate stops shows what one would expect from the greater phonetic duration of the consonantal occlusion alone. At least some singleton stops reveal articulatory undershoot of the target closure, as shown by the incomplete closures and tendency to spirantize in some dialects (compare discussion of tonal undershoot in Bruce 1990). With the pseudo-geminate stops having complete closures, and the longer closure duration providing more time for a pressure differential to build up, the pseudo-geminates stops are predicted to have stronger bursts than singleton stops. However, the difference in burst strength between singletons and pseudo-geminates was not very great if it was reliable at all. The real effect on burst strength resides in the comparison of singleton stops and true (or lexical) geminates. This is the crux of the argument that true geminates involve enhancement. This enhancement supports the contrastiveness of geminates in all structural positions, the unusual property of Berber phonology which motivated the study.

A question for future research is whether this enhancement is actually phonologically mediated, as Ridouane proposes by invoking rules on the feature [tense]. One alternative, suggested by work presented by Baese and Goldrick (2006) at the Paris meeting, is that contrastiveness directly affects speech production in a cascading activation model. Another alternative, suggested by the hybrid exemplar model of Pierrehumbert (2002), is that the memory traces of lexical geminates are somewhat hyperarticulated, because only hyperarticulated examples are reliably recognized in competition with their minimally different competitors. These two alternatives both predict differential effects on burst strength depending on whether the geminate is in a dense lexical neighborhood with singleton competitors or not. A third possibility, lying in between these proposals and Ridouane's is that the articulatory strengthening of the true geminates is a fully general effect which resides in the phonetic implementation rules. The antecedents for this analysis are work showing that prosodic position at the phrasal level can directly affect articulatory strength (Pierrehumbert and Talkin 1992; Keating, Cho, Fougeron, and Hsu 2003).

Like Byrd and Choi, Ridouane also observed individual variation in his recordings. In particular, the talkers in his study exhibited variation in their implementation of the secondary correlates of gemination, including devoicing, release RMS amplitude, and preceding vowel duration. As the role of secondary acoustic and articulatory correlates of phonological features are examined in greater detail, inter-talker variation in the use of those correlates should also be explored. Individuals may show significant variation in the cues that they use to signal contrast, particularly at the secondary level.

Together, the studies by Byrd and Choi, and Ridouane contribute to our understanding of the phonetic implementation of metrical structure in production. In both cases, the authors made some assumptions about the kinds of information that speakers and listeners "know," particularly with respect to the relationship between abstract linguistic units and continuous acoustic signals or the overlapping of articulatory gestures in real time. That is, the abstractions were assumed and their phonetic implementation was explored experimentally. But, can we assume that abstract levels of representation exist? Abstractions may be epiphenomenal, arising from the encoding and storage of millions of individual experiences. Or they may reflect external generalizations about a self-organizing system; as linguists grapple with the complexity of language, they may impart abstraction where it does not truly exist. Or, as is assumed by most language researchers, abstractions may be an inherent part of the cognitive system itself. We take up this issue in the next section.

4. Levels of representation

The nature of representations and the types of representations available for lexical and phonological processing in perception has been a continuous focus of laboratory phonology research over the years. In most cases, the researchers conclude by arguing for higher-level abstract representations of phonological, prosodic, and lexical structures. However, the extent to which these representations are fully specified phonologically and the nature of the matching process in recognition has been hotly debated (e.g., Lahiri and Marslen-Wilson 1992; Ohala and Ohala 1995; Newman, Sawusch, and Luce 2000).

The final paper in the 10th anniversary session, by Cutler et al., was specifically designed to explore the question of the status of abstract levels of representation. In a perceptual learning study, they obtained three distinct pieces of evidence from which they concluded that an abstract level of representation is required for some kinds of phonological processing. First, participants were able to learn a new acoustic /f/-/s/ boundary in a relatively short amount of time, despite years of experience with a different boundary. Second, the listeners generalized what they had learned to novel words; the new boundary affected their expectations about new cases. Third, the new boundary could be reactivated 12 hours after training. Learning, generalization, and retention over time all suggest abstraction over the training exemplars. If participants were only learning individual items, learning would be slow and laborious, generalization would not occur, and we would not expect retention over any considerable amount of time, particularly with intervening speech input.

Cutler et al. also crucially found that the learning was talker-specific, suggesting that the learners were operating in a social context; they only shifted their perceptual categories when the same voice was encountered, much as they would only shift their perceptual categories to match a specific foreign accent when talking to a speaker with that accent. We know that listening to tokens produced by multiple talkers aids in the learning of new phonological contrasts (e.g., Bradlow, Akahane-Yamada, Pisoni, and Tohkura 1999). However, multiple talkers can also impede processing in spoken word recognition tasks, resulting in slower and less accurate performance (Mullennix, Pisoni, and Martin 1989). As Cutler et al. note, several recent studies using a similar methodology have obtained mixed results with respect to the generalizability of learning across talkers. Thus, one question that deserves further investigation is: how do listeners know when to generalize to novel tokens produced by novel talkers and when to generalize only to novel tokens produced by the same talker?

We agree with Cutler et al.'s conclusion that abstract representations are real. In fact, we argue that abstract representations are necessary because the speech system is incredibly complex and highly dimensional. Twenty-two peripheral control elements are involved in controlling the larynx alone (Farley 1996) and Ridouane described five acoustic dimensions related to geminates alone. Systems with many degrees of freedom are hard to learn and hard to control. Trying to statistically optimize a high-dimensional space (i.e., learn the system) is difficult because high-dimensional optimizations get lost in hyperspace before settling on the optimal system. Trying to control a high-dimensional space (i.e., in production) leads to chaos because there are too many potential solutions to a given problem.

Fortunately, evolution solved the degrees of freedom problem for us by endowing us with neural circuits. Any neural implementation of logical or categorical structure is a circuit (or a network). These neural circuits reduce overall dimensionality in the system to help us deal with the high degrees of freedom in the real world. In addition, serial encoding means that neural circuits reduce active dimensions even further at any given point in time. Neural reduction of high-dimensional spaces is common throughout the animal kingdom for solving complex perceptual problems, such as rats' use of their whisker arrays to determine distance and terrain, and for object recognition (Schultz, Solomon, Peshkin, and Hartmann 2005). Thus, our neural circuitry provides the architecture for creating abstractions from the high-dimensional acoustic and articulatory spaces that we use for language (Beckman and Pierrehumbert 2003; Çetin and Ostendorf 2005).

More generally, variability in the environment breeds neural abstraction, which allows for flexibility and adaptive plasticity. As such, abstraction serves as the mechanism by which we attain plasticity in dealing with a highly variable environment. We have thus come full circle and are again faced with the inescapable truth that language, like all cognitive systems, is embodied. Our models of language and language use must therefore include an account of the relationship between mind and matter.

5. A new commodity

When talking about the linguistic system, the common conceptual division is one between the cognitive system, including segments and suprasegmentals (e.g., autosegmental-metrical structure), and the sociophonetic system, including talker identity, speaking style, dialect, etc. However, we would like to entertain an alternative division between small time scales (e.g., segments) and large time scales (e.g., suprasegmentals and language as a social system).

What we are proposing is a specific instantiation of our broader claims regarding variability, plasticity, and embodiment in which the critical division between the two sets of structures is one of dynamics. Vowel formant frequencies are eigenvalues of the vocal tract. Perception of these eigenvalues provides the listener with information about the vocal tract shape, which can then be used for vowel identification. In acquisition, articulatory exploration provides a correspondence between acoustic information and vocal tract configurations. Critically, however, these eigenvalues are achieved in a very short time window; one glottal pulse is adequate to produce the formant spectrum. As a result, physical models of vowel production and perception can be nearly stationary. However, as soon as you move to even the smallest prosodic structure (i.e., a CV syllable), the stationary approximation fails and dynamical models are required.

Thus, prosodic and social structures are linked at the larger time scale because they rely on dynamic models of the reallocation of attention and resources. The common mechanism tying these structures together is a network with loops. For prosodic structures, this network is made up of the neural circuits responsible for sequencing, grouping, and prominence. For social structures, the network is made up of the speakers in a population. Even the smallest production-perception loop conceived in *The Speech Chain* (Denes and Pinson 1963), required *two* brains, one for the talker and one for the listener. These networks are simultaneously active and, crucially, they interact. For example, prosody is one area in which you entrain to other people on both long time scales, as in acquisition, and short time scales, as in accommodation to your interlocutor's speaking rate or pitch range. Like the results reported by Cutler et al. for talker-specific learning, social information is also encoded differently depending on the task or the goals of the listener.

We challenge the laboratory phonology community to develop unified models of autosegmental-metrical structure and social structure. The precedents laboratory phonologists have set in autosegmental-metrical structure will be useful in this endeavor, as will exploration of mathematical parallels, such as dynamical network theory. Finally, the great historical successes in modeling the articulatory and acoustic foundations of segmental inventories provide the benchmark for the level of scientific explanation we hope to achieve in this new endeavor.

Acknowledgments

This work was supported by grants from the James S. McDonnell Foundation and NIH NIDCD (F32 DC007237) to Northwestern University. The authors would like to thank

Stefan Wuchty, a Northwestern Institute on Complex Systems postdoctoral fellow, for his assistance with the citation networks analysis.

References

Baese, Melissa and Matthew Goldrick 2006 Lexical effects on variation independent of phonotactics. Poster presented at Laboratory Phonology 10, Paris.

Beckman, Mary E. and Jan Edwards 1990 Lengthenings and shortenings and the nature of prosodic constituency. In: John Kingston and Mary E. Beckman (eds.), *Papers in Laboratory Phonology 1: Between the Grammar and Physics of Speech*, 152-178. Cambridge: Cambridge University Press.

Beckman, Mary, Jan Edwards, and Janet Fletcher 1992 Prosodic structure and tempo in a sonority model of articulatory dynamics. In: Gerard J. Doherty and D. Robert Ladd (eds.), *Gesture, Segment, Prosody* (Papers in Laboratory Phonology 2), 68-86. Cambridge: Cambridge University Press.

Beckman, Mary and Janet Pierrehumbert 2003 Interpreting "phonetic interpretation" over the lexicon. In: John Local, Richard Ogden and Rosalind Temple (eds.), *Phonetic Interpretation: Papers in Laboratory Phonology VI,* 13-37. Cambridge: Cambridge University Press.

Bradlow, Ann R., Reiko Akahane-Yamada, David B. Pisoni, and Yoh'ichi Tohkura 1999 Training Japanese listeners to identify English /r/ and /l/: Long-term retention of learning in perception and production. *Perception & Psychophysics,* 61: 977-985.

Broe, Michael 1992 An introduction to feature geometry. In: Gerard J. Doherty and D. Robert Ladd (eds.), *Gesture, Segment, Prosody* (Papers in Laboratory Phonology 2), 149-165. Cambridge: Cambridge University Press.

Bruce, Gösta 1990 Alignment and composition of tonal accents: Comments on Silverman and Pierrehumbert's paper. In: John Kingston and Mary E. Beckman (eds.), *Papers in Laboratory Phonology 1: Between the Grammar and Physics of Speech*, 107-114. Cambridge: Cambridge University Press.

Çetin, Özgür and Mari Ostendorf 2005 Multi-rate and variable-rate modeling of speech at phone and syllable time scales. *International Conference on Acoustics, Speech, and Signal Processing* I: 665-668.

Cho, Taehong and Patricia A. Keating 2001 Articulatory and acoustic studies on domain-initial strengthening in Korean. *Journal of Phonetics* 29: 155-190.

Denes, Peter B. and Elliot N. Pinson 1993 *The Speech Chain*, Second Edition. New York: W. H. Freeman.

Farley, Glenn R. 1996 A biomechanical laryngeal model of voice F0 and glottal width control. *Journal of the Acoustical Society of America* 100: 3794-3812.

Fougeron, Cécile and Patricia A. Keating 1997 Articulatory strengthening at edges of prosodic domains. *Journal of the Acoustical Society of America* 101: 3728-3740.

Harris, Katherine S. 1978 Vowel duration change and its underlying physiological mechanisms. *Language and Speech* 21: 354-361.

Hay, Jennifer 2003 *Causes and Consequences of Word Structure*. New York: Routledge.

Johnson, Keith, Edward Flemming, and Richard Wright 1993 The hyperspace effect: Phonetic targets Are hyperarticulated. *University of California Working Papers in Phonetics* 83: 55-73.

Keating, Patricia, Taehong Cho, Cécile Fougeron, and Chai-Shune Hsu 2003 Domain-initial articulatory strengthening in four languages. In: John Local, Richard Ogden and Rosalind Temple (eds.), *Phonetic Interpretation: Papers in Laboratory Phonology VI,* 145-163. Cambridge: Cambridge University Press.

Kraljic, Tanya and Arthur G. Samuel 2006 Generalization in perceptual learning for speech. *Psychonomic Bulletin and Review* 13: 262-268.

Ladd, D. Robert and James M. Scobbie 2003 External sandhi as gestural overlap? Counter-evidence from Sardinian. In: John Local, Richard Ogden and Rosalind Temple (eds.), *Phonetic Interpretation: Papers in Laboratory Phonology VI,* 164-182. Cambridge: Cambridge University Press.

Lahiri, Aditi and William Marslen-Wilson 1992 Lexical processing and phonological representation. In: Gerard J. Doherty and D. Robert Ladd (eds.), *Gesture, Segment, Prosody* (Papers in Laboratory Phonology 2), 229-254. Cambridge: Cambridge University Press.

McCarthy, John J. 1986 OCP Effects: Gemination and antigemination. *Linguistic Inquiry* 17: 207-263.

Mullennix, John W., David B. Pisoni, and Christopher S. Martin 1989 Some effects of talker variability on spoken word recognition. *Journal of the Acoustical Society of America* 85: 365-378.

Nespor, Marina, and Irene Vogel 1986 *Prosodic Phonology*. Dordrecht: Foris.

Newman, Rochelle, James Sawusch, and Paul Luce 2000 Underspecification and phoneme frequency in speech perception. In: Michael B. Broe and Janet B.

Pierrehumbert (eds.), *Papers in Laboratory Phonology V: Acquisition and the Lexicon*, 299-312. Cambridge: Cambridge University Press.

Ohala, John 1990 The phonetics and phonology of aspects of assimilation. In: John Kingston and Mary E. Beckman (eds.), *Papers in Laboratory Phonology 1: Between the Grammar and Physics of Speech*, 258-275. Cambridge: Cambridge University Press.

Ohala, John J. and Manjari Ohala 1995 Speech perception and lexical representation: The role of vowel nasalization in Hindi and English. In: Bruce Connell and Amalia Arvaniti (eds.), *Phonology and Phonetic Evidence: Papers in Laboratory Phonology IV*, 41-60. Cambridge: Cambridge University Press.

Pierrehumbert, Janet 1994 Prosodic effects on glottal allophones. In: O. Fujimura (ed.), *Vocal Fold Physiology 8*. San Diego: Singular Press.

Pierrehumbert, Janet 2002 Word-specific phonetics. In: Carlos Gussenhoven and Natasha Warner (eds.), *Laboratory Phonology 7*(Phonology and Phonetics 4), 101-140. Berlin: Mouton de Gruyter.

Pierrehumbert, Janet B. and Mary E. Beckman 1988 *Japanese Tone Structure* (Linguistic Inquiry Monograph Series No. 15). Cambridge, MA: MIT Press.

Pierrehumbert, Janet and D. Talkin 1992 Lenition of /h/ and glottal stop. In: Gerard J. Doherty and D. Robert Ladd (eds.), *Gesture, Segment, Prosody* (Papers in Laboratory Phonology 2), 90-117. Cambridge: Cambridge University Press.

Saltzman, Elliot, Anders Löfqvist, and Subhobrata Mitra 2000 'Glue' and 'clocks': Intergestural cohesion and global timing. In: Michael B. Broe and Janet B. Pierrehumbert (eds.), *Papers in Laboratory Phonology V: Acquisition and the Lexicon*, 88-101. Cambridge: Cambridge University Press.

Schultz, Aimee E., Joseph H. Solomon, Michael A. Peshkin, and Mitra J. Hartmann 2005 Multifunctional whisker arrays for distance detection, terrain mapping, and object feature extraction. *IEEE International Conference on Robotics and Automation*.

Smith, Caroline L. 1995 Prosodic patterns in the coordination of vowel and consonant gestures. In: Bruce Connell and Amalia Arvaniti (eds.), *Phonology and Phonetic Evidence: Papers in Laboratory Phonology IV*, 205-222. Cambridge: Cambridge University Press.

Xu, Yi 2004 Understanding tone from the perspective of production and perception. *Language and Linguistics* 5:757-797.