Cues or Contexts in Feature Licensing Constraints

Robert Kirchner, U.Ill./UCLA, November, 1996

0. Introduction

Steriade (1993, 1995, 1996) has proposed that the distribution of phonological contrasts is explained, not by a principle of Prosodic Licensing (Itô 1986, Goldsmith 1990) (licensing of features only in certain prosodic positions, e.g. onset position), but rather by the presence of sufficiently audible phonetic cues in the relevant contexts that signal these contrasts. For example, Steriade (1996) observes that contrastive glottalization in sonorants is frequently restricted to *post*-sonorant position (e.g. Yokuts), while in obstruents it is restricted to *pre*-sonorant position (e.g. Kashaya), where the relevant adjacent sonorant is voiced. In sonorants, glottalization is realized as creaky voice, typically beginning during the preceding segment; whereas in obstruents, it is realized as ejection, i.e. a loud release burst, followed by a long lag before the onset of voicing.[1] Let us assume that creaky voice is incompatible with obstruents, plausibly for aerodynamic reasons. Therefore, in post-obstruent or word-initial position (or following a voiceless sonorant, if such exist in the language), glottalization lacks the ancillary cue of (at least partially) creaky voicing during the preceding segment. For ejective obstruents, on the other hand, in word-final position, or preceding a voiceless segment, the cue of long voice onset time is absent; in pre-obstruent position, the release burst is obscured by the following stop closure or fricative noise; moreover, even if the following obstruent is voiced, the typically stronger "spontaneous voicing" in a sonorant makes for a clearer voice onset time cue in pre-sonorant than pre-obstruent position. This analysis can be elegantly formalized in terms of a set of constraints expressing the requirement of sufficient cues: "if [+constricted, +son] then [+partial creak] in the preceding segment"; and "if [+constricted, -son] then [+long voice onset time] and [+ejective burst]."[2] That is, we treat the contrastive feature and its licensing cues alike as phonological features, and formally express the dependencies directly between them. Call this Direct Licensing.

The problem, as Steriade observes, is that the presence or absence of a cue in a given token of some utterance is often not dispositive of the distribution of a licensed feature specification. Rather, it is whether the licensing cue is *typically* present or absent. For example, in word-initial position, a glottalized sonorant *may* be accompanied by anticipatory creak, if there is a preceding word which ends in a voiced sonorant; but this cue is not *reliably* present; therefore [+constricted glottis] does not appear in word-initial position, regardless of the preceding word. In such cases, Steriade assumes that the formalism must refer to a list of more abstract contexts in which the supporting cues are typically present: these contexts, not the cues themselves, are the formal licensers of the feature. Call this Indirect Licensing. In the glottalization case above, the constraints

---

[1] Cf. Silverman 1995 for an auditorily based explanation of these typical timing relations.

[2] Even more directly, assuming parallel auditory and articulatory representations à la Flemming 1995, we may abandon the feature [constricted glottis]: the constraints would then be "A [+creaky voice] domain must extend over more than one segment" and "if [+ejective burst] then [+long VOT]." The relation between the auditory features [creaky voice] and [ejective burst] lies in their common mapping to glottal constriction gestures in the articulatory representation.

would be "if [+constricted, +son] then /[+son,+voi]__" and "if [+constricted, -son] then /__[+son,+voi]." This Indirect Licensing approach is unsatisfying, however, since the formalism does not capture the relation between the actual licensing cues and the list of formal licensing contexts; or if it does, the procedure for computing this relation is not straightforward.

This paper will show that the Indirect Licensing approach is in fact unnecessary: the effect of feature stability in the face of token-to-token variation is an artefact of the role of faithfulness constraints in speech recognition and production, particularly in the face of sociolinguistic variation. Consequently, we can adopt the maximally elegant approach in the grammatical formalism: direct dependency between features and their licensing cues. We can, moreover, adopt a uniform representational treatment of potentially contrastive features and low-level phonetic cues.

1. Cue Variability in the Recognition Task.

1.1. The Inactivity of Licensing Constraints in the Recognition Task. Assume a grammar G s.t. the glottalized sonorant licensing constraint, "if [+constricted, +son] then preceding [+partial creak]" » PRES(constricted). Under G, a glottalized sonorant must be preceded by an (at least partially) creaky voiced segment. The non-contrastiveness of [partial creak] is guaranteed, either by low ranking of PRES(partial creak), or by assuming that there is no such constraint in UG (see Kirchner 1996). If [partial creak] is absent underlyingly, it will be filled in on the surface (1a). If in some context there is no compatible segment to carry the [partial creak] licensing cue, e.g. pre-obstruent position or word-initial position (in citation form), then [constricted glottis] cannot occur in this context either: rather, it neutralizes to [-constricted] (1b,c).

(1)

		*[-son, +partial creak]	LICENSING	PRES(constricted)
a. ☞	aɲ̃ - a̰a̰ɲ̃			
	aɲ̃ - aɲ̃		*!	
	aɲ̃ - an			*!
b.	agɲ̃a - agɲ̃a		*!	
	agɲ̃a - agg̃ɲ̃a	*!		
☞	agɲ̃a - agna			*
c.	ɲ̃a - ɲ̃a		*!	
☞	ɲ̃a - na			*

We will now see that even a listener equipped with a neutralizing grammar such as G can unproblematically recognize a lexical item containing some value of [constricted glottis], notwithstanding variation in the licensing cue, [+partial creak], in particular tokens. The crucial observation is that, since the licensing constraint (like OT well-formedness constraints generally) is a constraint *on surface representations*, it imposes no requirements on underlying representations, and therefore is irrelevant to the selection of the best UR for a certain surface form. That is, a given surface form, which is the input to the recognition task, will either violate or satisfy surface well-formedness constraints; but if the candidates are all mappings from that surface form to some UR, the

candidates will all tie w.r.t. the surface well-formedness violations, and so they are irrelevant to the selection:

(2)

		*[+F,+G]	PRES(F)	*[-F,-G]	PRES(G)
☞	+F,+G - /+F,+G/	*			
	+F,+G - /+F,-G/	*			*!
	+F,+G - /-F,+G/	*	*!		
	+F,+G - /-F,-G/	*	*!		*

Only constraints which refer to underlying representation, principally the input/output faithfulness constraints (Orgun 1995, Kirchner 1996), are potentially active in the recognition task.[3]

Assume a listener with grammar G, presented with an auditory input which lacks partial creak in the segment preceding a glottalized sonorant (3a), or which contains partial creak preceding a non-glottalized sonorant (3b) (this variation in the licensing cue can arise from any number of sources: background noise, a dialect discrepancy between listener and speaker, a sudden coughing fit on the part of the speaker, etc.)

(3)

			LICENSING	PRES(constricted)
a.	☞	an̰ - /an̰/	*	
		an̰ - /an/	*	*!
b.		a̰a̰ # na - /ba # n̰a/		*!
	☞	a̰a̰ # na - /ba # na/		

This tableau shows that, under G, the listener can correctly recognize the lexical item /an̰/, distinguishing it from /an/, despite the aberrant *absence* of [+partial creak]; and similarly for /na/ versus /n̰a/, despite the aberrant *presence* of [+partial creak].

1.2. The "Best Match" Relation. More generally, this shows that an OT grammar as applied to the recognition task simply looks for the "best match" between the input auditory representation and the set of URs. The constraint ranking (specifically, w.r.t. to the faithfulness constraints) establishes the priority of the cues in this best match relation.[4] This result may even provide, in rough outline, a new formal approach to the long-standing problem of inter-speaker variation in speech recognition (see Lindblom 1991); that is, how does one (generally, if not infallibly) recognize the lexical items intended by another speaker, notwithstanding considerable phonetic variation between

---

[3]Cf. Smolensky's (1996) related observation that a single OT grammar can give rise to both severely constrained production, and sophisticated recognition, as observed in certain stages of children's phonological development.

[4]The phenomenon of dialect adaptation (whereby one's ability to understand a speaker of a different dialect improves with exposure) can be handled as follows: as the listener establishes a number of matches between the speaker's surface outputs and the intended lexical items (aided, presumably, by top-down information, see below), the listener develops a speaker-specific ranking such that the regularly violated faithfulness constraints are demoted. This speaker-specific ranking then permits a more accurate mapping between new surface forms of this speaker and intended lexical items, hence improved recognition.

one's own realization of those items and the speaker's. For example, how can a female speaker of Mid-Atlantic American English, match the token [kœɑɯ] (with F1 at 710 Hz), produced by a male speaker of RP English, with her own lexical item /kœaˌ/ (with an F1 frequency typically around 820 Hz)? A plausible answer is that [kœɑɯ] is a closer match with /kœaˌ/ than with any other item in the listener's lexicon (e.g. [kœɔɯ] (*caw*)).

For this idea to work, however, we must say something non-trivial about the class of candidate URs. If the candidate set includes all possible URs, whether corresponding to existing lexical items or not, then *any* featural discrepancy between the surface token and existing lexical items will result in non-recognition. In the hypothetical above, for example, the listener would match [kœɑɯ] to /kœɑɯ/, for which she has no existing lexical entry, failing to match it with /kœaˌ/. On the other hand, if we restrict the set of candidates to existing lexical items, then an auditory input will *always* be matched with the closest lexical item, an equally undesirable result. For example an English listener presented with the input [bɔrʒfʎnaki] will unhesitatingly identify this as 'borscht phonology' (assuming this is the closest match) notwithstanding the large-scale phonetic differences between the matched forms (nor the semantic/pragmatic anomaly of the resulting expression). However, the non-recognition result can be obtained, even if the candidate set is restricted to existing lexical items, by including an additional candidate, identical to the input auditory representation (the "none-of-the-above candidate"), which always violates a constraint "*NON-RECOGNITION."

(4)

		PRES(F1• 700 Hz)[5]	PRES(stop closure)	*NON-RECOG	PRES(sharp F3 drop)[6]	etc.
☞	kœɑɯ - /kœaˌ/				*	
	kœɑɯ - /kœɔɯ/	*!				
	kœɑɯ - /kœɑɯ/ (none of the above)			*!		
	bɔrʒfʎnaki - /bɔˌʃt#fʎnalʎdʒi/		*!			
☞	bɔrʒfʎnaki - /bɔrʒfʎnaki/ (none of the above)			*		

The ranking of *NON-RECOGNITION relative to particular faithfulness constraints determines the extent to which a surface token can differ from an existing lexical item before the none-of-the-above candidate wins. (The more faithfulness constraints dominate *NON-RECOGNITION, the closer the resemblance between must be between a surface form and an existing lexical item.)

A further issue is the role of "top-down" syntactic, semantic, and pragmatic factors in this model of speech recognition. For example, given the utterance

---

[5]As an more sophisticated alternative to recognition of vowels based on closeness of raw formant frequency, cf. Bladon 1986, who suggests identification of vowels in terms of difference between formants, rescaled in auditory space.

[6]I.e. rhoticity.

[ðʌ pœɔjz ˌ hiˌ], the verb 'are' ([ˌ]) requires the preceding subject to be plural. But how can we avoid an exact match with the existing lexical item /pœɔjz/ ('poise') in favor of the worse match /bɔj+ z/ ('boy-pl.'), or non-recognition (i.e. "What's a poy?") (inspired by an anecdote of Hockett?). The answer is to allow the requirement of syntactic agreement to knock out 'poise' as a candidate. This can be done (a) by allowing the candidate-generating function GEN to generate only top-down-appropriate lexical items as candidates (in this case, only plural forms, thereby excluding 'poise'); (b) by simultaneous evaluating syntactic well-formedness and closeness of phonetic match in a single OT tableau (if the syntactic AGREEMENT constraint dominates the phonological constraints, all candidates with singular nouns such as 'poise' will lose relative to plural forms /bɔj+z/ or (unrecognized stem) /pœɔj+z/); or (c) by allowing the phonology to select 'poise' as the first-round winner, which is then found to be unparsible in the syntactic component, whereupon the listener backtracks to phonological evaluation, this time excluding the output of the previous (syntactically unsuccessful) attempt, 'poise', and so on. By the same mechanism, we can block recognition of lexical items which are semantically or pragmatically anomalous. Thus, top-down knowledge (combined with relatively high ranking of *NON-RECOGNITION) can force the identification of lexical items, even in the face of massive signal degradation (e.g. casual speech in a noisy room) (cf. Lindblom 1991).

2. Cue Variability and the Production Task

In light of the irrelevance of licensing constraints in speech recognition, how then can we say that a ranking such as LICENSING » PRES(constricted) has any effect on the distribution of [constricted glottis]? To answer this, we must consider the much greater role of feature licensing in the *production* task. Given grammar G (1), it is impossible for a surface phonological representation to contain [+constricted, +son] without a preceding segment containing [+partial creak], and absence of the licensing cue in even a single surface representation will force neutralization to [-constricted glottis] in that representation. This result reflects the observation that speakers can handle a great deal more variation in the speech that they recognize than they themselves produce.

2.1. Cue Variation Due to Unforeseen Production Factors. Nevertheless, variable realization of licensing cues may arise, without neutralization, even in production, in two ways. First, it must be borne in mind that the surface phonological representation is a mental state, reflecting the speaker's *plan* for the utterance; however, assuming the possibility of error in the translation of this mental plan into the corresponding articulatory event (and its acoustic consequences), the licensing cue may occur in less than 100% of the tokens in which the licensed feature occurs. For example, the planned utterance [tatitʰaa̰n̰a] might be realized as [ta-*cough*-n̰a]. That is, the speaker's plan assumes a normal, unirritated trachea, but this "best-laid plan ... gangs awry," due to unforeseen production factors.

2.2. Featural Stability Due to Difficulty of Learning Typically Neutralized Specifications. Furthermore, consider the word-initial context. Under G, a word-initial underlyingly glottalized sonorant will neutralize to [-constricted glottis] in all syntactic

contexts except following a word-final voiced sonorant, which can carry the [+partial creak] licensing cue.

(5) /n̰a/ - [na] / d # __ /n̰a/ - [n̰a] / a̰a̰ # __
- [na] / k # __
- [na] / a̰ # __

Let us assume, however, that a learner, in setting up URs for lexical items, posits the feature specifications which are *typically* present in the surface tokens (a variant on the notion of Lexicon Optimization, Prince and Smolensky 1993). The learner would therefore posit /na/ rather than /n̰a/. Once we assume that the [-constricted glottis] specification is present in the UR, the presence or absence of [+partial creak] has no effect on the surface value.

(6)

		LICENSING	PRES(constricted)
	na - n̰a		*!
☞	na - na		
	a # na - a̰a̰ # n̰a		*!
☞	a # na - a # na		

Therefore, in word-initial contexts, we predict stably non-glottalized sonorants, notwithstanding the possible presence of the licensing cue in particular tokens.

This story, at first blush, seems to rule out all manner of morphophonemic alternations. For example, if in Russian the typical tokens of some nominal stem are unaffixed (nominative singular) forms, e.g. [ras] ('time'), [les] ('forest'), and these are taken as the URs for these stems (/ras/, /les/), how could Russian have a voicing distinction which shows up in the affixed (e.g. dative) forms of the stems, [razu] vs. [lesu]? The answer is that the dative forms are not necessarily derived "on-line" from the nominative: they may have their own lexical entries (/razu/, /lesu/). Nor does this assumption mean that we give up on the observation that the [ras]/[razu] and [les]/[lesu] paradigms show great phonological identity between the base and the affixed form. Using output/output faithfulness constraints (independently motivated by over- and under-application effects in reduplication, and cyclicity effects, see McCarthy and Prince 1995, Benua 1995, Flemming 1995, Kenstowicz 1995), we can enforce phonological identity (to a greater or lesser degree) between morphologically related surface forms, without deriving the surface forms from a common UR.

(7)

		PRES(cont, Base/ Derivative)	FINAL DEVOICING	PRES(voi, UR/PR)	PRES(voi, Base/ Derivative)	PRES(cont, UR/PR)
☞	/ras/ - ras /razu/ - razu				*	*
	/ras/ - ras /razu/ - rasu			*!		

	/ras/ - raz /razu/ - razu	*!	*	
☞	/rat/ - ras /razu/ - razu			* *
	/rat/ - rat /razu/ - razu	*!		

Thus, the paradigm [rat]/[razu] is impossible under this ranking, since the surface forms differ w.r.t. [continuant], though [ras]/[razu] is possibile, since the surface forms merely differ w.r.t. lower-ranked PRES(voi,Base/Derivative). This does not mean that surface [ras] actually derives from /rat/ (since /ras/ is a better input); rather, this tableau shows that learner could not assign [rat] and [razu] to the same paradigm under this ranking.

2.3. Featural Stability Due to Sociolinguistic Variation in the Activity of the Licensing Constraint. It is possible for members of the same general speech community to have subtly but significantly different grammars, an observation frequently invoked in theories of diachrony. Assume that some members of the speech community have grammar G (the neutralizers), while others have grammar G', in which PRES(constricted) » LICENSING (the non-neutralizers). Further assume an undominated constraint, in both G and G', which forces the absence of the licensing cue in some narrow context, e.g. "*[+partial creak] in a nuclear stressed vowel]." For the neutralizers, then, glottalized sonorants neutralize following a nuclear stressed vowel.

(8)

	G	* ːV̰ LICENSING	PRES(constricted)
	an̰ - 'an̰ (in stressed position)	*!	
	an̰ - 'a̰an̰ (in stressed position)	*!	
☞	an̰ - 'an (in stressed position)		*
	an̰ - an̰ (elsewhere)	*!	
☞	an̰ - a̰an̰ (elsewhere)		
	an̰ - an (elsewhere)		*!

For the non-neutralizers, however, glottalized sonorants can occur in any context, regardless of the absence of the licensing cue, [+partial creak]; even word-initially.

(9)

	G'	* ːV̰ PRES(constricted)	LICENSING
☞	an̰ - 'an̰ (in stressed position)		*
	an̰ - 'a̰an̰ (in stressed position)	*!	
	an̰ - 'an (in stressed position)	*!	
	an̰ - an̰ (elsewhere)		*!
☞	an̰ - a̰an̰ (elsewhere)		
	an̰ - an (elsewhere)	*!	
☞	n̰a - n̰a (in either position)		
	n̰a - na (in either position)	*!	

The interesting observation is that the non-neutralizers may systematically lack lexical items such as /n̰a/, while having items such as /an̰/ (contrasting with /an/), despite the well-formedness of [n̰a] under their grammar.[7] Such a gap can be explained in terms of the influence of the neutralizers on the shared vocabulary[8] of the speech community, particularly if the neutralizers are more numerous, are perceived as a high prestige group, or have other sociolinguistic factors favoring the perpetuation of their pronunciations. The same is true if a generation of neutralizers diachronically preceeds a generation of non-neutralizers. That is, a non-neutralizer learning vocabulary items principally from neutralizers (directly or indirectly) will never encounter items containing a word-initial glottalized sonorant. Moreover, the non-neutralizers will able learn the distinction between /an̰/ and /an/, since the *typical* realization of /an̰/ by the neutralizers is [a̰an̰], inducing the non-neutralizers to set up a UR containing a glottalized sonorant.

Now, the problem of indirect dependence in production between features and the cues that license them is solved. Rather than grammaticize the indirect dependence in terms of lists of contexts in which cues are typically present, as in the Indirect Licensing approach, we can capture the effect of indirect dependence in terms of sociolinguistic interaction. In other words, in the non-neutralizers' dialect, [+constricted glottis] is licensed by the occurrence of [+partial creak] *not in their own PRs, but in the typical PRs of the neutralizers, from whom they learn their vocabulary.*

Of course, this state of affairs is diachronically unstable: it breaks down as soon the neutralizers' linguistic domination of the non-neutralizers ends, and the neutralizers begin to learn vocabulary items such as [n̰a] from other speech communities. This interaction, incidentally, captures the standardly observed diachronic progression from a "phonetic" rule (neutralization of feature conditioned by the actual presence of some cues in a given token, among the true neutralizers) to its "grammaticized" reflex, a phonological rule (conditioned by some more abstract context, among the non-neutralizers), which then acquires lexical exceptions (as the non-neutralizers achieve sociolinguistic independence).

3. Epilogue: the Recognition Task Revisited

It is natural to think of the licensing relation between features and cues (in the Direct Licensing approach, between features and other features) in terms of the recognition task. That is, requirements of multiple cues serve to enhance the auditory robustness of contrasts, making them more easily perceived by the recognizer. Counterintuitively, however, we have discovered that licensing constraints play a role in production, but not in recognition. We will now demonstrate that the *effect* of such licensing contraints on production does in fact substantially aid in the task of speech recognition, as we would expect.

---

[7]This assumes that the absence of [n̰a] in the data does not itself force a learner to adopt the neutralizing ranking.

[8]I use the "vocabulary" and "vocabulary items" as the E-language correspondents of "lexicon" and "lexical items," which I treat as I-language notions.

Consider the task of recognizing an auditory input [na], with a lexicon containing /na/, /n̰a/, and /la/.

(9)

	Input: [na]	Pres(nasal)	Pres(constricted glottis)	etc.
☞	/na/			
	/n̰a/		*!	
	/la/	*!		

The listener will recognize the input as a token of /na/, since this appears to be a perfect match. However, if the speaker intended /n̰a/, but the percept of glottalization was obliterated by background noise, a misidentification will result. Indeed, the difficulty of reliably perceiving glottalization in the absence of the licensing cue is precisely what the LICENSING constraint expresses.

Now assume a speech community which homogeneously neutralizes [constricted glottis] in the absence of [+partial creak], as in tableaux (1) and (8) above. These speakers produce no instances of word-initial glottalized sonorants; they therefore encounter no tokens containing word-initial glottalized sonorants. Consequently, there is no item corresponding to /n̰a/ in the speakers' lexicons. The ambiguity in the recognition task is accordingly reduced.

(10)

	Input: [na]	Pres(nasal)	Pres(constricted glottis)	etc.
☞	/na/			
	/la/	*!		

By constraining production, the speech community indirectly constrains the range of URs that have to be considered in the recognition task, specifically eliminating URs with distinctive glottalization in environments where that glottalization is difficult to perceive.

APPENDIX

Using these insights about recognition in an OT grammar, the following result falls out: high frequency lexical items, such as affixes (especially inflectional) and function words will neutralize before content words. *This accounts for the observation that affixes and function words typically have impoverished contrasts compared to roots and content words, without stipulating the result in terms of distinct faithfulness constraints for roots vs. affixes or content vs. function words.*

Assume some language with casual speech neutralization of ʃ to s. FREQ is a constraint expressing the quicker access of high-frequency lexical items.

	Auditory input: [lekis]	FREQ	*NON-RECOG	PRES(ant)
☞	/lek-iʃ/ (‘talk-past’)			*

/lekis/ (‘pancreas’)	*!		
lekis (n.o.t.a.)	?	*!	
/lekiʃ/ (‘gunwale’)	*!		*

Now, that we’ve matched a token [-is] to the ‘past’ suffix /-iʃ/, assume that this somehow alters the UR for this suffix. In production, what are you faithful to now? The [-ant] original specification, or the [+ant] specification of the new token? Assume that as a preponderance of [+ant] tokens accumulates, we are no longer faithful to [-ant] at all. As a result, the past-suffix has neutralized to /-is/. Now only [+ant] stridents are possible in suffixes.

	Auditory input: [lekiʃ]	(some syntactic constraint that knocks out verbs in this position)	FREQ	*NON-RECOG	PRES(ant)
	/lek-iʃ/ (‘talk-past’)	*!			*
	/lekis/ (‘pancreas’)		*		
	lekis (n.o.t.a.)		?	*!	
☞	/lekiʃ/ (‘gunwale’)		*		*

REFERENCES

Benua, L. (1995) Identity Effects in Morphological Truncation, *UMass Occasional Papers in Linguistics* 18, 77-136.

Bladon, A. (1986) "Phonetics for Hearers," in G. MacGregor, ed., *Language for Hearers*, Oxford: Pergamon Press, 1-24.

Flemming, E. (1995) *Auditory Features in Phonology*, Doctoral Dissertation, UCLA.

Goldsmith, J. (1990) *Autosegmental and Metrical Phonology*, Oxford: Blackwell.

Itô, J. (1986) *Syllable Theory in Prosodic Phonology,* Doctoral Dissertation, University of Massachusetts, Amherst.

Kenstowicz, M. (1995) *Base-Identity and Uniform Exponence*, Ms., MIT.

Kirchner, R. (1996) *Faithfulness and Contrastiveness,* Ms., UCLA.

McCarthy, J. and A. Prince (1995) *Faithfulness and Reduplicative Identity,* Ms., University of Massachusetts, Amherst, and Rutgers University, New Brunswick, New Jersey.

Orgun, O. (1995) *Correspondence and Identity Constraints in Two-Level Optimality Theory,* Ms., U.C.-Berkeley (available on the Rutgers Optimality Archive at ruccs.rutgers.edu).

Prince, A. and P. Smolensky (1993) *Optimality Theory: Constraint Interaction in Generative Grammar,* Ms., Rutgers University, University of Colorado.

Silverman, D. (1995) *Phasing and Recoverability*, Ph.D. Dissertation, UCLA.

Smolensky, P. (1996) On the Comprehension/Production Dilemma in Child Language, Ms., Johns Hopkins University.

Steriade, D. (1993) Positional Neutralization, paper presented at NELS.

Steriade, D. (1995) *Neutralization and the Expression of Contrast,* Ms., UCLA.

Steriade, D. (1996) *Licensing Laryngeal Features,* UCLA Working Papers in Phonology 2.