Neighborhood Density and the Root-Affix Distinction

Adam Ussishkin and Andrew Wedel

University of Arizona and UC Santa Cruz

1.   Introduction

Work in phonology and morphology consistently recognizes that affixes show a strong tendency to be less marked than roots crosslinguistically. Within Optimality Theory, this observation has been dealt with through the imposition of a universally-fixed ranking, known as the Root-Affix Faithfulness Metaconstraint (McCarthy and Prince 1995), which states that faithfulness to root material universally outranks faithfulness to affix material.

In this paper, we present evidence drawn from psycholinguistic research that provides functional grounding for the root-affix markedness distinction without resorting to any stipulative device or metaconstraint. Research into lexical access processes shows that the frequency and lexical neighborhood characteristics of a target lexical entry have strong effects on its efficiency of access. The Neighborhood Activation Model, which accounts for these findings by modeling access as a process of lexical competition, predicts that because affixes are significantly more frequent than roots, they will tend to contain less marked material. This prediction stems from the proposal that it is *effective contrast*, a function of both frequency and phonemic contrast, that is optimized relative to markedness, not phonemic contrast alone. Thus, despite the relatively low phonemic contrast of material in affixes crosslinguistically, their high frequency allows them efficient processing and access, leading them to evolve toward less marked structure.

Further, we show that the Neighborhood Activation Model predicts that when language-specific factors result in different affix neighborhood relations, such as in Hebrew, the Root-Affix Faithfulness Metaconstraint may be reversed. Hence, the ostensibly anomalous behavior of affixes in languages like Hebrew can be straightforwardly accommodated by abandoning the metaconstraint strategy in favor of a model of contrast based on psycholinguistic evidence that the efficiency of lexical access

is dependent on the interaction between lexical neighborhood characteristics and frequency.

2.   Markedness in Roots vs. Affixes

Crosslinguistically, affixes make use of less marked elements than roots, whether segmental or structural. For example, affixes make greater use of less marked, e.g. coronal segments, and tend to avoid marked segments, e.g. pharyngeals. Affixal material shows a preference for peripheral vowels and the use of core syllable types as well. The following examples illustrate several cases in which roots contain material that segmentally more marked than material found in affixes.

(1)   English

     In English, affixal consonants are overwhelmingly coronal; this restriction is found in many other languages as well.

(2)   Salish

     In Salish, glottalized consonants are only found in roots and lexical suffixes; grammatical suffixes do not contain glottalized consonants (S. Urbanczyk, p.c.).

(3)   Frisian and Dutch

     Frisian and Dutch affixes utilize a less marked subset of the vowel inventory available to roots (W. Visser, p.c.).

(4)   Turkish

     Turkish roots contain mid, rounded vowels, but affixes do not.

     In addition to the observation that roots tend to contain a more marked segmental inventory than affixes, roots also tend to contain a more marked structural inventory than affixes. In particular, roots tend to exhibit a superset of prosodic shapes available to affixes. The following examples illustrate this distinction.

(5)   Tibetan

     In Tibetan, syllabic affix morphemes exhibit only simple onsets, though complex onsets are widely attested in roots (C. Beckwith, p.c.).

(6)   Sanskrit

     In Sanskrit, roots contain complex onsets, but affixes contain only simple onsets (McCarthy and Prince 1995).

3.   Optimality Theory and the Root-Affix Markedness Distinction

Within the framework of Optimality Theory (henceforth OT; Prince and Smolensky 1993), the markedness distinction between roots and affixes has been acknowledged via the "Root-Affix Faithfulness Metaconstraint" (McCarthy and Prince 1995):

(7)   Root-Affix Faithfulness Metaconstraint

       FAITH-ROOT » FAITH-AFFIX

     The Root-Affix Faithfulness Metaconstraint (henceforth, RAFM) stipulates that faithfulness to root material is unexceptionally higher-ranking than faithfulness to affix material. As a consequence, roots are allowed to violate some markedness constraint(s) that affixes may never violate, with the result that affixes contain a subset of the elements permitted to occur in roots. Within OT this is formalized by the ranking FAITH-ROOT » C » FAITH-AFFIX, where C is some markedness constraint (McCarthy and Prince 1995).

     As an example, consider a language such as Arabic that allows pharyngeal consonants in roots but never in affixes. The markedness constraint against pharyngeal consonants, *PHAR, is outranked by FAITH-ROOT, since pharyngeals occur in root material. Hence, any underlying pharyngeal consonants will be faithfully realized if they occur within a root. However, the markedness constraint *PHAR is higher-ranking than FAITH-AFFIX, since pharyngeals are never found in affixes. This latter ranking explains why affixes never contain pharyngeal consonants; an underlying pharyngeal in an affix cannot faithfully surface. The situation is formalized through the ranking FAITH-ROOT » *PHAR » FAITH-AFFIX.

     This approach works well to explain the observations described above, but also makes a clear prediction: namely, that there exist no languages in which faithfulness to affixes may outrank faithfulness to roots. Additionally, the metaconstraint does not provide any explanation for why the ranking it instantiates should be important, nor does it give any principled reason for the difference in behavior between roots and affixes.

     Furthermore, counterexamples exist to the prediction that no language may require greater faithfulness to affixes than to roots; e.g., ATR harmony in Turkana (Noske 2000) and stem-mutating suffixes in Slavic. Additionally, Ussishkin (1999, 2000) finds that Hebrew, because of strict prosodic requirements limiting the shape of verbs, requires a grammar in which root faithfulness is outranked by affix faithfulness. The very existence of such cases, therefore, undermines the metaconstraint-based approach, which by its stipulative nature does not easily accommodate exceptions.

     The major goal of this paper is to point the way to a more deeply articulated, functionally grounded explanation for the root-affix distinction. This approach accounts for the overwhelming tendency for affixes to be less marked than roots through reference to well-studied effects in lexical access phenomena, and in doing so predicts the conditions under which this tendency may be violated.

4.   Neighborhood Density

Over the last decade, research on the process of lexical access has shown that the efficiency with which a lexical entry is recognized is affected by the number of other phonemically similar entries in the lexicon. In lexical decision tasks, for example, the speed and accuracy with which an experimental subject can decide if a presented word is a real word or not is inversely related to the number of similar words in the subject's lexicon (Goldinger, Luce, and Pisoni 1989, Cluff and Luce 1990, Luce and Pisoni 1998).

Similarity has been operationally defined for most studies of this effect by the 'add, subtract or change' rule, in which a lexical entry counts as similar to another if it can be changed into the other by adding, subtracting or changing one phoneme. For example, under this definition entries that are similar to *cat* include *scat*, *at*, *bat*, *cot* and *cap*. The set of lexical entries that are similar to a given entry by this measure are termed its lexical *neighbors*, and a *lexical neighborhood* is such a set of minimally phonemically contrastive lexical entries[1].

Returning to the lexical decision task, experiments show that subjects are able to decide that a word with few neighbors like *orange* is a real word more rapidly and accurately than a word with many neighbors like *cat*. This is termed the neighborhood density effect, and is illustrated graphically below in (8).

(8)   Neighborhood density and access efficiency

Low density access efficiency > *High density access efficiency*
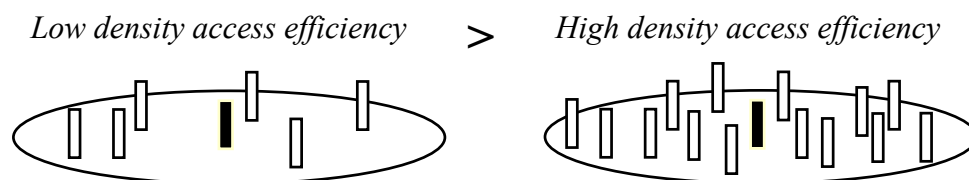
In (8), ovals represent lexical neighborhoods and bars represent lexical entries. The distance between bars represents the degree of phonemic similarity. The lexical entry represented by the solid bar on the left is accessed more efficiently because it has relatively few neighbors. On the right, the lexical entry represented by the solid bar has many more near neighbors, and is found to be accessed less efficiently.

4.1.   The Neighborhood Activation Model

The neighborhood density effect has been accounted for in the Neighborhood Activation Model (NAM) of Luce and Pisoni (1998; cf. also references therein) which proposes that a stimulus input (for instance, a spoken word) activates entries in the lexicon proportionally to their degree of similarity to the stimulus. Activated lexical entries then compete for selection as the intended match to the stimulus. Final selection of a lexical entry is made in part then on the basis of differential levels of activation, with the lexical entry with greatest relative activation most likely to be selected.

---

[1] This operational definition is of course crude – it does not take into account degree of perceptual similarity between sounds, for example – but the neighborhood density effect is apparently robust enough to be detected even with this relatively unsophisticated measure.

A stimulus matching an entry in a dense neighborhood will broadly activate many lexical items to a similar degree, resulting in a smaller *relative* difference in activation between the best fitting lexical item and its nearest competitors. The closer competition between these minimally contrastive lexical entries under NAM accounts for the observation that lexical entries in dense neighborhoods are accessed with reduced efficiency. Over the last decade, a considerable array of experiments using a variety of methodologies has supported the predictions of NAM (Goldinger, Luce and Pisoni 1989, Cluff and Luce 1990, Metsala 1997, Newman, Sawusch and Luce 1997, Luce and Pisoni 1998, Vitevitch and Luce 1998, Boyczuk and Baum 1999, Bradlow and Pisoni 1999, Vitevitch et al. 1999, Dirks et al. 2001).

4.2. Frequency and Contrast

Alongside neighborhood density, lexical frequency also has a strong effect on the efficiency of lexical access: all else being equal, the higher a lexical item's token frequency, the more rapid and accurate its access (cf. Dirks et al. 2001). In fact, it has been found that a high token frequency can mitigate the deleterious effects of a dense neighborhood on lexical access. If a target lexical entry has a higher token frequency than its neighbors, access is efficient even if the neighborhood is dense. This is illustrated below in (9).

(9)   High relative frequency and efficiency of access

Low density access efficiency ≈ *High density access efficiency*
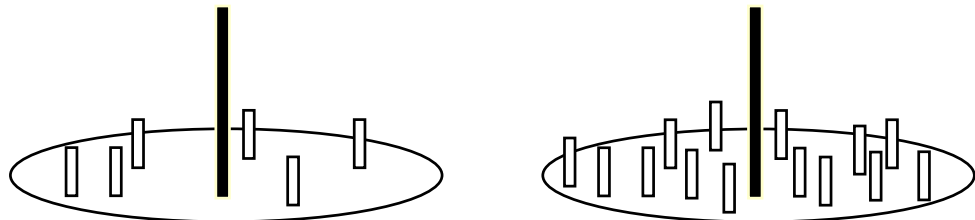
In (9), ovals represent lexical neighborhoods; bars represent lexical entries. The distance between bars represents degree of phonemic similarity, and the height of the bars represents the token frequency. On the left, the lexical entry represented by the solid bar is accessed relatively efficiently because it is very frequent relative to its lexical neighbors, which are, in addition, few and far between. On the right, the lexical entry represented by the solid bars has many more neighbors, but is still accessed efficiently because it is much more frequent than any of those neighbors.

The diagrams in (9) illustrate the finding that a high relative token frequency renders access efficiency relatively insensitive to neighborhood density; i.e., a high frequency lexical item in a dense neighborhood may be accessed nearly as efficiently as

one in a low density neighborhood. Thus, efficiency of access is affected by (at least[2]) two factors. On the one hand, the lower a lexical entry's phonemic contrast with other entries (i.e., the denser its neighborhood), the lower its efficiency of access. On the other hand, the more often a lexical entry is accessed (i.e., the higher its token frequency), the more efficient that access is.

In terms of lexical access then, the success with which a lexical item contrasts with its neighbors is a function both of its phonemic contrast and its relative token frequency. To refer to this composite property of lexical entries, we coin the term *effective contrast*. Because effective contrast is sensitive to both phonemic contrast and relative token frequency, two lexical items can have the same effective contrast but be in differently dense neighborhoods. For example, an infrequent lexical item that is in a sparse neighborhood may have the same effective contrast as a relatively frequent one in a dense neighborhood[3].

Effective contrast as conceived here is directly related to lexical access. In functional terms, lexical access may be the most significant point for contrast: lexical items that are accessed efficiently are by definition those that contrast well. We choose the term effective contrast to make clear both the relationship and the distinction between phonemic contrast and contrast as reflected in access efficiency.

5.   Markedness and Effective Contrast

5.1.   Effective Contrast and the Lexicon

Despite a crosslinguistic tendency to avoid phonologically marked elements, it is a commonplace that every language makes use of some set of marked elements in order to provide sufficient contrast between lexical items. But what does 'sufficient' contrast mean in this context? If it refers to the ability of a given utterance to be reliably and efficiently mapped to a unique lexical entry, then the influence of frequency on lexical access renders the notion of contrast measured solely in terms of phonemes inadequate. Because frequent lexical items are accessed more efficiently than infrequent ones in dense neighborhoods, language processes that function to support phonemic contrast should be particularly important for less frequent lexical items. Conversely, as markedness reduction inevitably also results in loss of phonemic contrast, this should be best tolerated in the most frequently accessed lexical entries. In the balance between maintaining contrast and minimizing markedness then, we propose that it is *effective* contrast that is optimized with respect to markedness, not phonemic contrast.

5.2.   The Effective Contrast of Affixes Relative to Roots

Affixes are distinct from roots in at least two different ways relevant to the current discussion. First, like all function morphemes, affixes have a much higher token

---

[2] The access efficiency of a lexical item has also been found to depend to a lesser degree (Dirks et al. 2001) on the aggregate relative frequency of its near neighbors. This value is termed the *neighborhood frequency* (Luce 1986, Luce et al. 1990, Luce and Pisoni 1998).

[3] This difference corresponds to the 'hard word' vs. 'easy word' distinction used in Bradlow and Pisoni (1999).

frequency than roots, that is, an individual affix is likely to occur much more frequently than any given root (Segalowitz and Lane 2000 and therein). Second, as function morphemes, affixes as a group have a much lower *type* frequency than roots, that is, there are vastly fewer affixes than there are roots.

There are two general models of the lexical storage of affixes and other function morphemes relative to content morphemes. In one model (termed here the 'single lexicon model'), content and function morphemes are stored together in a single lexicon (cf. Segalowitz and Lane 2000). Under this model, all lexical entries, whether of function or content morphemes, compete with each other in lexical access. In a second general model (termed here the 'split lexicon model'), function and content morphemes are stored in distinct lexicons, and so cannot compete with one another in lexical access (cf. Bradley 1978, Biassou et al. 1997)[4].

Under the single lexicon model in which affixes and roots are stored together, the much greater token frequency of affixes relative to roots means that their access will be efficient even if their phonemic contrast is low – in other words, their effective contrast remains high because of the compensatory effects of high frequency. Under the split lexicon model in which affixes are stored separately from roots, even though any given affix entry is not likely to be accessed more frequently than its affix neighbors, the low absolute number of affixes means that the affix lexicon will necessarily be sparsely populated. Under this model, we expect affixes to be efficiently accessed even if phonemic contrast is low, because any given lexical entry has so few near neighbors. Under either model then, the effective contrast of affixes is high, whether due to their high token-, or low type-frequency.

Above we proposed that in the balance between reduction of markedness and maintenance of contrast, it is effective contrast that is relevant. This understanding of the tug-of-war between maintenance of contrast and minimization of markedness predicts that reduction in markedness will have the least impact for morphemes with the highest effective contrast, whether due to a sparse neighborhood or high relative token frequency. As effective contrast is already very high for affixes, they should, as a group, tolerate greater reduction in markedness than word-classes with lower average effective contrast, such as roots, where an equivalent reduction in markedness would have a greater impact on effective contrast.

6.   Theoretical Consequences

In Hebrew, most derivational affixes occur as patterns of two vowels, resulting in the familiar system of verbal classes or binyanim. This type of system has been analyzed as a classic case of root and pattern morphology (McCarthy 1979, 1981), and within OT a different view has recently been proposed by Ussishkin (1999, 2000). Under this recent approach, Hebrew requires a ranking in which faithfulness to affixes outranks all other faithfulness constraints. This obviously contradicts the RAFM, and is necessary in order to achieve what is known as melodic overwriting. Hebrew verbs are prosodically restricted to two syllables, and are formed by concatenating a bisyllabic base form with a bivocalic affix. Given the upper limit of two syllables, the resulting complex form can

---

[4] For example, in our previous example listing some lexical neighbors of *cat*, under the split lexicon model, *scat, bat, cot*, and *cap* are neighbors, but *at* is not.

only accommodate the affixal material by deleting base material[5]. The following section details the relevant data.

6.1.   The Hebrew Pattern

As an example of this process, consider the following paradigm, which presents verbs derived from the base form *gadal* 'he grew'.

(10)   The verbal paradigm for *gadal*, 'to grow'

Base form	*+affix*	*Derived form*	*Gloss*
gadal			'he grew'
	i e	gidel	'he raised'
	u a	gudal	'he was raised'
	hi i	higdil	'he enlarged'
	hu a	hugdal	'he was enlarged'

In each case, it is clear that the derived form is in a sense missing part of the original base form: namely, the vowels of the base. It should also be clear that this is the only way in which the bivocalic affix can be expressed, so in such a system a ranking that contradicts the RAFM should be expected to occur. In other words, if this metaconstraint did actually hold universally a language like Hebrew would be predicted as impossible. Given the existence of Hebrew, therefore, the RAFM may only be retained as a universal tendency and not as a universal absolute.

6.2.   NAM Predicts the Hebrew Pattern

Under the single lexicon model, NAM predicts that affixes will tolerate markedness reduction better than roots because they are more frequent. Likewise, under the split lexicon model, the low number of affixes means that affixes will be unlikely to have many near neighbors, again resulting in a high tolerance for markedness reduction. However, any special conditions resulting in lower effective contrast for affixes, such as an unusually high number of high frequency neighbors, should make affix faithfulness a higher priority.

In Hebrew, affixes for different verbal classes are composed of two vowels. Given the five-vowel inventory of Hebrew (*i, e, a, o, u*) a total of 25 possible affixes exist. Nineteen of these are actually instantiated, with the result that phonemic contrast between most verbal affixes cannot be further reduced without neutralization to homophony. Therefore, because of special restrictions on what may serve as verbal affixal material, verbal affixes are all near neighbors of one another. This is illustrated below in (11) under the single lexicon model[6].

---

[5] For a comprehensive analysis of the nature of the two-syllable size limit imposed on Hebrew verbs, see Ussishkin (2000).

[6] Under the split lexicon model, (11) would be identical but for the absence of the unfilled bars, with the same consequences as those described here for the single lexicon model.

(11)   Affix neighborhood density in Hebrew

Most languages   *Hebrew*
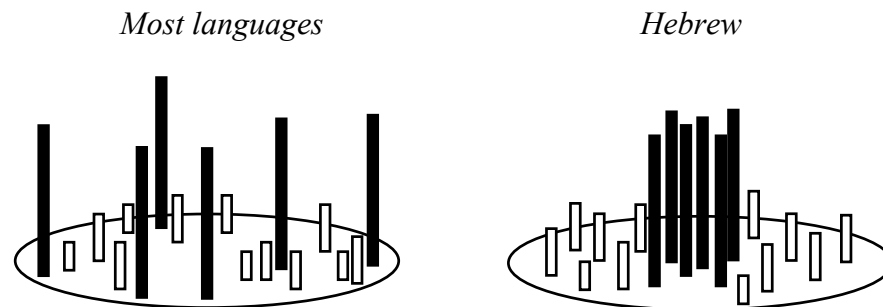
In (11), the ovals represent regions of the lexicon and the bars represent lexical entries. Bar height represents token frequency. The solid bars are representative of affix entries. For the general case depicted on the left, affixes are distributed relatively evenly throughout lexical space, due to the facts that that affixes are relatively few and that affixal material is drawn from a relatively large set of elements. For Hebrew, represented on the right, affixes are taken from a very limited set of similar elements, composed of vowel-vowel sequences drawn from a set of five vowels. Therefore, Hebrew verb class affixes are grouped very closely together in lexical space.

Above we saw that the high frequency of affixes relative to their neighbors makes their access efficient even when phonemic contrast is low. In Hebrew however, we see that verbal affixes are *not* more frequent than their near neighbors, because those neighbors are affixes as well. Unlike in other languages then, the high frequency of a verbal affix in Hebrew cannot compensate for a low phonemic contrast, with the result that phonemic contrast must be maintained to preserve adequate effective contrast. Hebrew demonstrates that affix faithfulness must outrank root faithfulness in cases where following the RAFM would obliterate the contrast between affixes entirely[7].

7.   Conclusion

In the tradition of grounding phonological patterns in functional constraints, we have illuminated a possible link between low markedness in affixes and recent advances in experimental linguistics showing that both frequency and neighborhood density interact in lexical access.

In OT, 'contrast' is the ability of lexical items to be differentiated from one another. The functional correlate of contrast in this sense may be lexical access efficiency, which measures the degree to which lexical items are distinctive in their lexical neighborhood. Recent psycholinguistic research shows however that phonemic

---

[7] In cases where an affix is composed solely of a floating feature or a single segment, the constraint REALIZE-MORPHEME (e.g., Samek-Lodovici 1993, Rose 1997, Walker 1998, Kurisu 2001) may be invoked to ensure parsing of the morpheme. But why does REALIZE-MORPHEME only protect against loss of *all* contrastive material? As presented above, effective contrast is a function of both frequency and phonemic contrast. When phonemic contrast of an item approaches zero (a common fate of affixes), effective contrast begins to approach zero as well, irrespective of high frequency. This may explain why REALIZE-MORPHEME requires some exponence of a morpheme, even if such exponence is minimal.

contrast — that is, the simple difference between lexical items in terms of phonemes — is not the sole component of lexical access efficiency. Lexical neighborhood characteristics such as density and relative frequency have significant effects on access, with the result that a high frequency lexical item can be accessed efficiently even if its phonemic contrast is low.

Faithfulness constraints protect against erosion of contrast incurred by satisfaction of markedness constraints. In the balance between maintenance of contrast and reduction of markedness, we propose that it is some function of both frequency and phonemic contrast, here termed effective contrast, that is optimized in individual lexical entries rather than phonemic contrast alone. The idea that individual lexical entries require individual faithfulness constraints has been proposed in earlier work; for instance, see Ito and Mester's (1999) work on lexical stratal effects.

As discussed here, affixes are frequent relative to roots, and can therefore maintain a high effective contrast despite low phonemic contrast even if they are stored alongside roots in a single lexicon. If affixes are stored separately from roots, the fact that they are much less numerous directly results in a higher default effective contrast. Under either model, affixes therefore pay a lower penalty in effective contrast for reductions in markedness than roots, leading to the conclusion that effective contrast/markedness optimization should result in lower relative affix markedness.

McCarthy and Prince's Root-Affix Faithfulness Metaconstraint can therefore be grounded in the distinct requirements for phonemic contrastiveness placed on roots versus affixes. Languages in which the metaconstraint does not appear to hold, such as Hebrew, are predicted to present special circumstances that mitigate this distinction.

References

Biassou, N., Obler, L. K., Nespoulous, J.-L., Dordain, M., and Haris, K. S. 1997. Dual processing of open- and closed-class words. *Brain and Language* **57**:360-373.

Boyczuk, B. and Baum, S. 1999. The influence of neighborhood density on phonetic categorization in aphasia. *Brain and Language* **67**:46-70.

Bradley, D. 1978. *Computational Distinctions of Vocabulary Types*. Ph.D. Dissertation, MIT.

Bradlow, A., and Pisoni, D. 1999. Recognition of spoken words by native and non-native listeners: Talker-, listener-, and item-related factors. *Journal of the Acoustical Society of America* **106**:2074-2085.

Cluff, M. and Luce, P. 1990. Similarity neighborhoods of spoken two-syllable words: Retroactive effects on multiple activation. *Journal of Experimental Psychology: Human Perception and Performance* **16**:551-563.

Dirks, D., Takayanagi, S., Moshfegh, A., Noffsinger, P. D., and Fausti, S. 2001. Examination of the Neighborhood Activation Theory in normal and hearing-impaired listeners. *Ear and Hearing* **22**:1-13.

Goldinger, S., Luce, P., and Pisoni, D. 1989. Priming lexical neighbors of spoken words: Effects of competition and inhibition. *Journal of Memory and Language* **28**:501-518.

Ito, J., and Mester, A. 1999. The phonological lexicon. In N. Tsujimura, ed., *A Handbook of Japanese Linguistics*. Cambridge: Blackwell.

Kurisu, K. 2001. *The Phonology of Morpheme Realization*. Ph.D. dissertation, UC Santa Cruz.

Luce, P. 1986. A computational analysis of uniqueness points in auditory word recognition. *Perception and Psychophysics* **39**:155-158.

Luce, P., Pisoni, D., and Goldinger, S. 1990. Similarity neighborhoods of spoken words. In G. Altman (ed.), *Cognitive models of speech perception: Psycholinguistic and computational perspectives* (pp. 122-147). Cambridge, MA: MIT Press.

Luce, P., and Pisoni, D. 1998. Recognizing spoken words: The Neighborhood Activation Model. *Ear and Hearing* **19**:1-38.

McCarthy, J., and Prince, A. 1995. Faithfulness and reduplicative identity. In Beckman, J. N., Dickey, L. W., and Urbanczyk, S., eds., *University of Massachusetts Occasional Papers* **18**: *Papers in Optimality Theory*, 249-384.

Metsala, J. 1997. An examination of word frequency and neighborhood density in the development of spoken-word recognition. *Memory and Cognition* **25**:47-56.

Newman, R., Sawusch, J. and Luce, P. 1997. Lexical neighborhood effects in phonetic processing. *Journal of Experimental Psychology* **23**:873-889.

Noske, M. 2000. ATR harmony in Turkana: A case of Faith-Suffix » Faith-Root. *Natural Language and Linguistic Theory* **18**:771-812.

Prince, A., and Smolensky, P. 1993. Optimality Theory: Constraint interaction in generative grammar. Ms., Rutgers University and University of Colorado, Boulder.

Rose, S. 1997. *Theoretical Issues in Comparative Ethio-Semitic Phonology and Morphology*. Ph.D. dissertation, McGill University.

Samek-Lodovici, V. 1993. Morphological gemination. Talk given at ROW-1, October, 1993.

Segalowitz, S., and Lane, K. 2000. Lexical access of function versus content words. *Brain and Language* **75**:376-389.

Ussishkin, A. 2000. *The Emergence of Fixed Prosody*. Ph.D. Dissertation, UC Santa Cruz.

Vitevitch, M., Luce, P. 1998. When words compete: levels of processing in perception of spoken words. *Psychological Science* **9**:325-329.

Vitevitch, M., Luce, P., Pisoni, D. and Auer, E. 1999. Phonotactics, neighborhood activation and lexical access for spoken words. *Brain and Language* **68**:306-311.

Walker, R. 1998. *Nasalization, Neutral Segments, and Opacity Effects*. Ph.D. dissertation, UC Santa Cruz. [published by Garland, New York, 2000]

Department of Linguistics
University of Arizona
Douglass Building
1100 E. University Ave
Tucson, AZ 85716

ussishki@email.arizona.edu

Adam Ussishkin and Andrew Wedel

Department of Linguistics
Stevenson College
UC Santa Cruz
Santa Cruz, CA 95064

wedel@ling.ucsc.edu