Vowel Height: Reconsidering Distinctive Features
Author(s): Don Salting
Proceedings of the Twenty-Fourth Annual Meeting of the Berkeley Linguistics Society: General Session and Parasession on Phonetics and Phonological Universals (1998), pp. 391-402

VOWEL HEIGHT: Reconsidering Distinctive Features

Don Salting
Communication Disorders Technology, Inc., Indiana University

1. Introduction

A model of vowel height consisting of a two-tiered, symmetrical hierarchy of autonomous nodes is presented as a descriptor of the segmental organization for languages which exhibit [ATR] type harmonies. This model is called the Nested Subregister model, and is illustrated in (1) below as it would describe a typical nine-vowel inventory.

(1) The Nested Subregister Model

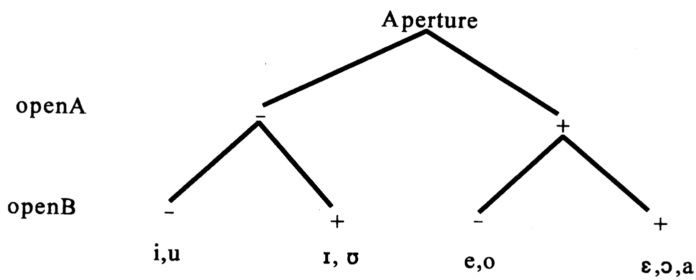

In the Nested Subregister model, the feature [openA] divides the height dimension in half. The [-openA] half is the less open, and thus the higher half, and the [+openA] half is the more open, and thus, the lower half. The feature [openB] represents a subregister, or subdivision of [openA]. As we will see in the two languages examined in this paper, the segmental makeup of the terminal nodes can vary, determined by the vowel inventory of the specific language. Following Clements (1991), a phonetic constant is that the leftmost (least open) vowels will always be the highest in the inventory, and the rightmost (most open) will be the lowest, with the remaining vowels arrayed by relative height.

The notion of an inventory-driven division of vowel features is contrary to the notions regarding distinctive features put forth in SPE. The traditional features as put forth in SPE are articulatorily based. They reference raising or lowering of the tongue body in relation to a "neutral" position defined as that in the English word 'bed' (Chomsky and Halle 1968:304). The assumption is that all languages divide the vowel space along the same parameters of musculature. It may be that the need for cooccurrence constraints and cleanup rules in so many harmony analyses stems from this assumption. In contrast, Archangeli & Pulleyblank (1994:135) cite evidence for cross-linguistic variability in the phonetic realization of F-elements. The Nested Subregister model allows for this sort of phonetic variability within the framework of a highly constrained hierarchy.

VOWEL HEIGHT: Reconsidering Distinctive Features

Don Salting
Communication Disorders Technology, Inc., Indiana University

1. Introduction

A model of vowel height consisting of a two-tiered, symmetrical hierarchy of autonomous nodes is presented as a descriptor of the segmental organization for languages which exhibit [ATR] type harmonies. This model is called the Nested Subregister model, and is illustrated in (1) below as it would describe a typical nine-vowel inventory.

(1) The Nested Subregister Model

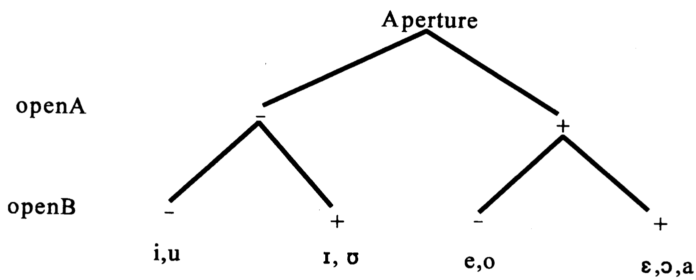

In the Nested Subregister model, the feature [openA] divides the height dimension in half. The [-openA] half is the less open, and thus the higher half, and the [+openA] half is the more open, and thus, the lower half. The feature [openB] represents a subregister, or subdivision of [openA]. As we will see in the two languages examined in this paper, the segmental makeup of the terminal nodes can vary, determined by the vowel inventory of the specific language. Following Clements (1991), a phonetic constant is that the leftmost (least open) vowels will always be the highest in the inventory, and the rightmost (most open) will be the lowest, with the remaining vowels arrayed by relative height.

The notion of an inventory-driven division of vowel features is contrary to the notions regarding distinctive features put forth in SPE. The traditional features as put forth in SPE are articulatorily based. They reference raising or lowering of the tongue body in relation to a "neutral" position defined as that in the English word 'bed' (Chomsky and Halle 1968:304). The assumption is that all languages divide the vowel space along the same parameters of musculature. It may be that the need for cooccurrence constraints and cleanup rules in so many harmony analyses stems from this assumption. In contrast, Archangeli & Pulleyblank (1994:135) cite evidence for cross-linguistic variability in the phonetic realization of F-elements. The Nested Subregister model allows for this sort of phonetic variability within the framework of a highly constrained hierarchy.

The departure from the divisions of vowel height referenced by traditional features is illustrated in Figs. (2) and (3) below.

(2) Vowel Space Divided by Traditional Features

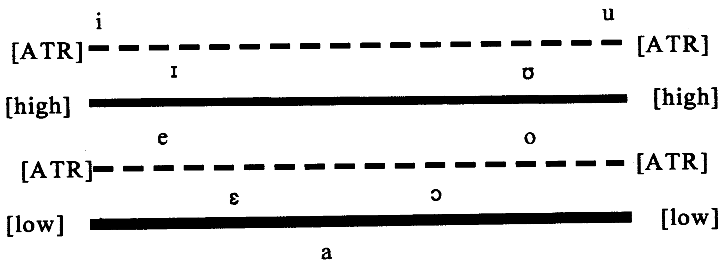

Notice in (2) that the feature [low] separates the lowest segment from the rest of the inventory. The feature [high] then divides the remainder, with [ATR] subdividing the two divisions created by [high].

(3) Vowel Space Divided by Nested Subregister Model

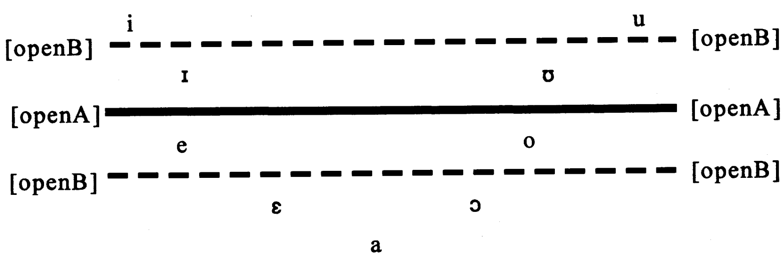

In (3), as mentioned above, the feature [openA] divides the entire vowel-height space, while [openB] subdivides the two registers created by [openA]. Adopting this arrangement to account for a harmony process, one could just as easily argue for eliminating [low] from the language in question. One could then account for the divisions in (3) as spaces defined by [high] and [ATR] as they are currently understood. Likewise, one could consider /ɛ,ɔ/ to be [+low]. But, as will be seen, the segmental inventory of each subregister will vary cross-linguistically. Thus, phonetic correlates as would be referenced by traditional features cannot be viewed as fixed for the nodes in the Nested Subregister model.

Linguistic parallels for the subdivided hierarchy in the Nested Subregister model can be seen in models for tone put forth in Yip (1980), Clements (1981) and Pulleyblank (1986).

After a brief background, two languages with very different harmony systems will be examined which are both accounted for equally well by the Nested Subregister model (hereafter NS model). Rather than a cursory look at a number of systems, this paper will examine two languages in relative detail and offer what is hopefully an explicit analysis. Further examples can be seen in Salting (1998).

2. Theoretical Background

The Aperture node and the feature [open] are adopted from Clements (1991). Also, following the notion of a feature hierarchy as put forth in Clements (1991), the nodes in the NS model are autonomous. Thus, in terms of feature spreading, segments from the inventory given in (1) above would be represented as in (4) below:

(4) Segmental Representation in Nested Subregister Model

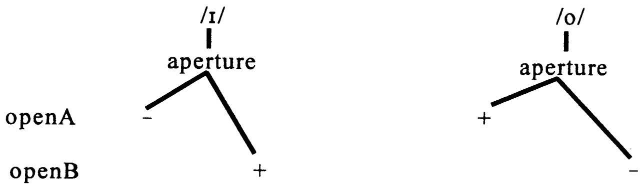

The NS model departs from Clements (1991) on several points, however. First, the NS model employs letters (A,B) to denote successive tiers rather than numbers. Clements (1991) introduces successively numbered degrees of [open] ([open1, open2, etc.]) to account for Bantu languages which exhibit scalar harmony systems (Nzebi, Kimatuumbi, i.a.). It is felt that Clements (1991) accounts very well for scalar harmony systems and these will not be addressed in this paper. The NS model is seen as a typological complement to Clements' scalar system.

For Bantu languages which exhibit non-scalar harmonies (Kinande, Sesotho), Clements (1991) introduces the notion of a branching hierarchy of the feature [open]. However, in a branching hierarchy, such as the NS model, successive tiers do not adhere to the definition of [open] as "...a uniform phonetic and phonological parameter..." (Clements 1991:38). Rather, the tiers denote subsets. Clements (1991) recognizes that the use of numbers for the hierarchic nodes is misleading, but argues for their use based on the fact that the languages in question (Bantu) are all closely related. In Salting (1998), data from two dialects of Italian (Southern Umbro and Northern Salentino) are offered to illustrate that this sort of inter-familial disparity occurs elsewhere. Therefore, in the NS model I adopt the use of a more arbitrary method of denoting hierarchic tiers, and use letters ([openA, openB]).

Second, in hierarchies of Clements (1991), the tier defined by [open1] is analogous to the feature [low], [open2] is analogous to [high], and [open3] is analogous to [ATR]. In the NS model, no reference is made to the traditional features of vowel height.

I assume several cardinal tenets of feature geometry:

(i) "Spreading is the sole mechanism of assimilation." (Hayes 1986:467)

(ii) The Constituent Spreading Hypothesis: "A single rule may spread no more than a single node of the feature hierarchy." (Pulleyblank 1988:314)

3. Lhasa Tibetan

Lhasa Tibetan has a 12 vowel inventory: /i,ü,u,ɪ,ɨ,ʊ,e,ö,o,ɛ,ɔ,a/. In addition to a full complement of vowels, Lhasa exhibits raising, lowering, and fronting harmony, which provide multiple triggers and targets in one language. These must, at some level, coexist. The data in this section are taken from Chang & Shefts (1964), Chang & Chang (1968), Dawson (1980) and Nornang (1978).

Raising harmony occurs across consonants and can affect any non-high vowel. This process is seen when the future suffix *-ki* is attached to verb roots:

(5) Raising Harmony

	BASE	FUTURE	GLOSS
(a)	šee	šii-ki	'get'
(b)	kho	khu-ki	'hear'
(c)	tsõõ	tsũ-ki	'sell'
(d)	phöö	phüü-ki	'flee'
(e)	nɛɛ	nɪɪ-ki	'sleep'
(f)	lɔɔ	lʊʊ-ki	'read'
(g)	šaa	šɨɨ-ki	'leave'

This process also occurs across word boundaries in compound nouns. Raising harmony in Lhasa is not problematic for traditional features, and can easily be analyzed as the simple spreading of [+high] within the domain of the phonological word. But, as we will see, the other harmony processes in Lhasa prove problematic when analyzed with standard vowel-height features.

Both lowering and fronting harmonies occur in the morphophonology of Lhasa. When a root ends in a single, oral vowel, certain suffixes will create a harmonized, word-final long vowel. If the harmony environment is not met, the morpheme will surface as a default CV suffix. Some data are given below. For the sake of space, only harmonizing forms are given.

(6) Lowering Harmony

	BASE	HARMONIZED	GLOSS	SUFFIX
(a)	ci	cɪɪ	'heavy'	Comparative
(b)	tshĩĩpu	tshĩĩpʊʊ	'stitch(tailor)'	Agent/Resident
(c)	qhare	qharɛɛ	'kind(s)'	Dative/Locative
(d)	soso	sosɔɔ	'oneself'	Dative/Locative

(7) Fronting Harmony

	BASE	HARMONIZED	GLOSS	SUFFIX
(a)	chu	chüü	'water'	Genitive
(b)	šiŋpɨ	šiŋpɪɪ	'farmer'	Genitive
(c)	cha	chɛɛ	'tea'	Instrumental

Salting (1995) and Salting (1998) argue that Lhasa has two allomorphs for each of the harmonizing suffixes, and that the harmonizing form consists of a mora with a single feature attached. In traditional terms, the active feature for the lowering harmony is [-ATR] and that for the fronting harmony is [-back]. For this paper, I assume the vowel place geometry of Odden (1991) which consists of a Back/Round node dominating the features [back] and [round].

Notice in (6c) that when /e/ is lowered, the result is [ɛ], as one would expect with [-ATR] spread. Notice, however, in (7c), that when /a/ is fronted, it also surfaces as [ɛ]. Assuming the traditional features, the fronting of /a/ would create a disallowable combination of *[-back, +low], forcing a cleanup rule to create [-low]. An option would be to say that /ɛ,ɔ,a/ are all [+low] in Lhasa, but then lowering harmony to /e/ would create a surface [ɛ] specified for *[-low]. This would require an ad hoc rule to change [-low] to [+low]. In either case, traditional features cannot represent Lhasa vowels such that all harmony processes are representative of a single internal organization of vowel features.

Yet another alternative would be to eliminate the notion of the feature [low] altogether. This is accomplished with the NS model. As with the option above, this situation implies that /ɛ,ɔ,a/ are all the same height, but now they are not [+low]. Rather, they are all members of the lowest subregister. They are distinguished from each other solely by their Back/Round status. The vowels for Lhasa would be arrayed as in the hierarchy in (8):

(8) Lhasa Vowel Hierarchy: the Nested Subregister Model

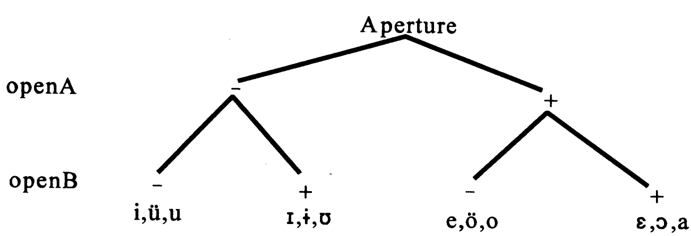

Given the above configuration, the two height-harmonies in Lhasa can be expressed with one simple rule each, and no clean-up conditions are required. For the raising harmony, the rule is *spread [-openA]*. A derivation of the future tense for 'sleep' from (5e) is given in (9) below. For the sake of space, all irrelevant tiers are omitted from this and all subsequent derivations.

(9) Raising Harmony in Lhasa

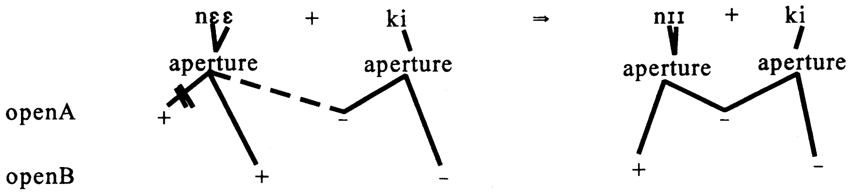

More telling however, are the derivations in (10) and (11) in which /e/ and /a/ both surface as [ε] respectively due to occurrences of lowering and fronting.

In (10) below, we see an example of the lowering harmony pattern. Following Salting (1995, 1998), the harmonizing suffix consists of a mora with a feature attached, in this case, [+openB]. The token is from (6c).

(10) Lowering Harmony in Lhasa

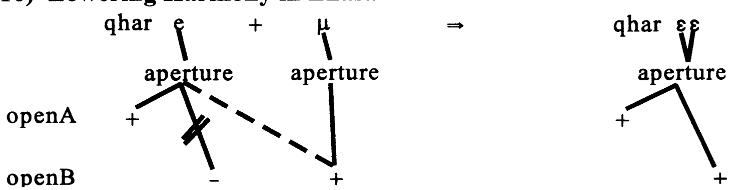

In (11), we see [ε] surfacing as a result of harmony with its other harmonic counterpart, /a/. This is due to fronting harmony. Again, we assume that the underlying trigger morpheme is a mora with a feature attached, this time, [-back]. The token in (11) is from (7c).

(11) Fronting Harmony in Lhasa

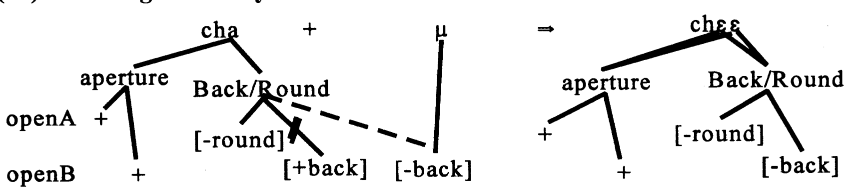

Most significant in (10) and (11) is the fact that the height features require no secondary rules or conditions to account for disallowable feature combinations. With the NS model, all harmony patterns in Lhasa fall out as natural reflections of an underlying organization within the phonological component. All segmental changes in Lhasa are shown in (12) below, with underlying and surface forms arrayed in respective order.

(12) Lhasa Harmonies

				Rule
(a)	Raising	e,o,ö,ɛ,ɔ,a	→ i,u,ü,ɪ,ʊ,ɨ	*spread [-openA]*
(b)	Lowering	i,u,e,o	→ ɪ,ʊ,ɛ,ɔ	*spread [+openB]*
(c)	Fronting	u,o,a	→ ü,ö,ɛ	*spread [-back]*

An argument could easily be made against the above analysis by claiming that either [low] doesn't exist in Lhasa, or that it remains unspecified until very late in the phonology. However, unrelated languages with differing inventories and processes provide evidence that the NS model represents a universal organizational parameter for vowel height. One such language is Ogori.

4. Ogori

Ogori is an Eastern Kwa language from Nigeria. It exhibits vowel-harmony patterns that are very difficult to analyze with traditional features, but it can be analyzed with the NS model without recourse to fill-in rules or constraints within the height component. All data are from Chumbow (1982). Ogori has a seven vowel inventory: /i,u,e,o,ɛ,ɔ,a/. There is a word-level constraint in Ogori barring the cooccurrence of /e,o/ with /ɛ,ɔ/. The vowels /i,u,a/ can occur with vowels from either of the 'mid' sets. Instances of vowel harmony in Ogori occur in the allomorphy of verb prefixes, in personal pronouns and in reduplication. For the first two cases, two distinct sets of vowels and a uniform pattern fall out. Reduplication exhibits a different pattern of vowel subsets than the other two types of harmony. I will illustrate these first two harmony processes then offer an analysis using hierarchical subsets. After that, reduplication will be examined as independent evidence for the analysis.

In the prefixal and pronominal harmonies in Ogori, two harmonic classes of vowels appear, which Chumbow (1982) labels as [+ATR] and [-ATR]. Each [+ATR] vowel is shown with its harmonic [-ATR] counterpart in (13) below:

(13) Ogori Harmonic Pairs

[+ATR]	[-ATR]
i	ɛ
u	ɔ
e	a
o	ɔ

Evidence for this pairing can be seen in the allomorphs for Ogori personal pronouns. These are given in (14) below:

(14) Ogori Personal Pronouns

	SUBJECT		OBJECT	
	Sg.	Pl.	Sg.	Pl.
1p	ì/ɛ̀	tì/tɛ̀	mú/mɔ́	tú/tɔ́
2p	ù/ɔ̀	nì/nɛ̀	ú/ɔ	nú/nɔ́
3p	è/à	bì/bɛ̀	é/á	bé/bá

When the verb contains any one of the [-ATR] vowels, the pronoun will agree in its specification for that feature as is shown in (15):

(15) Ogori Pronoun Harmony

1st Person Singular		2nd Person Singular		3rd Person Singular	
ì-jé	'I call'	ù-jé	'you call'	è-jé	'he calls'
ɛ̀-nɛ́	'I fling'	ɔ̀-nɛ́	'you fling'	à-nɛ́	'he flings'
ì-roró	'I think'	ù-roró	'you think'	è-roró	'he thinks'
ɛ̀-kpɔ	'I climb'	ɔ̀-kpɔ	'you climb'	à-kpɔ	'he climbs'

Evidence for the o ~ ɔ natural class is seen in (16) with the Agentive prefix:

(16) Ogori Agentives

	Verb	Gloss	Agentive	Gloss
a.	sìútu	'do a job'	ò-sìútu	'worker'
b.	bírépe	'plant crops'	ò-birépe	'planter'
c.	mɔ̀ɔtɛ́lɛ́	'make a pot'	ɔ̀-mɔ́ɔ̀tɛ́lɛ́	'potter'
d.	jájɔ	'dance'	ɔ̀-jájɔ	'dancer'

The active feature of the harmony can be seen in the Habitual prefix *dèkì-\dàkɛ̀-*. Though this prefix harmonizes, it blocks harmony to the word-initial pronoun, thus giving evidence that the default state is [+ATR].

(17) Harmony Blocking Affixes

bí dàkɛ̀ bɛ́ bù úmú → (after truncation) [bí dàkɛ̀ bɛ́ búmú]
they HAB beat their goat 'They always beat their goat'

With the data in (17) we can assume that, employing traditional features, the harmony rule for Ogori is *spread [-ATR]*. However, with one exception, all harmonic pairs require a variety of additional cleanup rules when analyzed with traditional features. This is illustrated in (18):

(18) Ogori Feature Changes

	[+ATR]	→	[-ATR]	ADDITIONAL FEATURE CHANGES
a.	i	→	ɛ	[+high] → [-high]
b.	e	→	a	[-low] → [+low]
c.	u	→	ɔ	[+high] → [-high]
d.	o	→	ɔ	no change

The most incongruous is (18b). Chumbow (1982:77) argues that it occurs "...in order to insure optimal utilization of all available vowels in the harmony process..." Were that the case, one would expect to find more cases of /e/ lowering to [a] 'over' [ɛ] cross-linguistically, but this seems to be very rare. On the other hand, one might equally expect /o/ to lower to [a] as well, instead of lowering to [ɔ] as it does. Clearly, the traditional features of vowel height do not account for there being anything natural about the natural classes in Ogori.

However, when analyzed with the NS model, harmony in Ogori can be explained with one rule and no conditions or constraints. As with Lhasa, the NS model divides the inventory along a center line and then subdivides each major division. Notice in (19) however, that the 'center' line as defined by [openA] separates different groups of segments for Ogori than it does for Lhasa (c.f. (8) above).

(19) Ogori Hierarchy

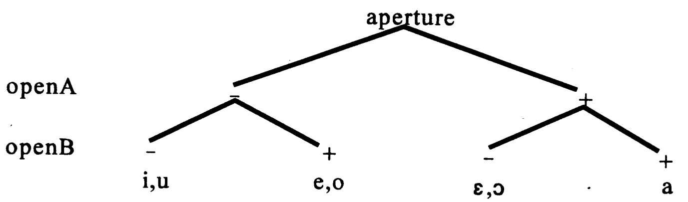

With this configuration, the harmony rule for Ogori is *spread [+openA]*. Some derivations from the data in (15) are given in (20,21) below.

(20) /ì + nɛ́/ → [ɛ̀nɛ́] *'I fling'*

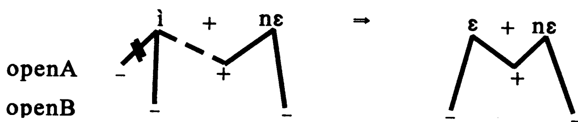

In the derivation in (21), we get a straightforward account of /e/ -> [a] in Ogori harmony.

(21) /è + kpɔ/ → [àkpɔ] *'he climbs'*

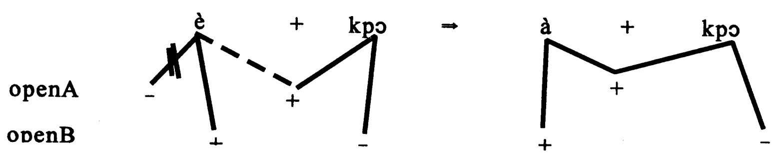

In both (20) and (21), the NS model is able to account for Ogori harmony in a way that does not require extra rules or constraints; all that is required is a simple rule spreading one feature. Furthermore, the NS model provides a paradigm that explains and predicts the natural classes. There is one exception, however, which does require a constraint. This occurs in /o/ → [ɔ] (Agentives) where the surface form also requires a change in the specification for [openB]. Though a change is required in the height component, it is required, not because the featural combination is disallowed of itself, but because it cannot cooccur with the underlying Back/Round specification. Ogori has no segment which is *[+openA, +openB, +round]. Therefore, the Back/Round status of the vowel takes precedence over its height; the restriction is not against a specific height, but against a specific height-roundness combination. The issue of Place-Height interface will surface in the next section on reduplication as well.

A good corroboration for the NS model would be found in a separate harmony pattern in the same language -- a pattern which exhibits a different harmonic grouping. Such a pattern is found in Ogori reduplication. Reduplication in Ogori occurs in the nominalization of verbs and adjectives. The initial syllable of the root is reduplicated and the nominalizing prefix ò/ɔ̀ which we saw in the Agentives (see (16) above) is added.

(22) Ogori Reduplication

A. VERBAL NOUNS

	Root	Verbal Noun	Gloss
1.	džé	ò-dzí-dze	'eating'
2.	rúwá	ò-rù-rúwá	'division'
3.	wá	ɔ̀-wɛ́-wa	'drinking'
4.	šá	ɔ̀-šɛ́-ša	'coming'
5.	bàlɛ́	ɔ̀-bɛ́-balɛ	'watching'

B. ADJECTIVAL NOUNS

	Root	Noun	Gloss
1.	gbódí	ò-gbí-gbodi	'bigness'
2.	kéré	ò-kí-kere	'smallness'
3.	lɔ́rɛ́	ɔ̀-lɛ́-lɔrɛ	'longness'
4.	dɔ́	ɔ̀-dɛ̀-dɔ́	'oldness'
5.	bɛ́rɛ́	ɔ̀-bɛ́-bɛ́rɛ	'softness'

With the exception of *rúwá*, all reduplicative affixes contain either [i] or [ɛ]. Within the framework of the NS model, this means that the reduplicative affix is an open syllable containing a front vowel specified [-openB] which gets its [openA] status from the root vowel. The harmonic pairings of Ogori reduplication are presented in (23) below:

(23) Ogori Reduplication Schematized

	RED. AFFIX	ROOT VOWEL
a.	u	u
b.	i	e
c.	i	o
d.	ɛ	a
e.	ɛ	ɛ
f.	ɛ	ɔ

As with the previous harmony processes, one can assume a default status of [-openA] for the reduplicative affix making the underlying vowel /i/. With this, the rule for reduplication is the same as for the other processes, namely *spread [+openA]*. A sample derivation is given in (24) below. The subsequent change of [openB] for the nominalizing prefix is as per the discussion above.

(24) /o-wi-wa/ ⇒ [ɔwɛwa] 'drinking'

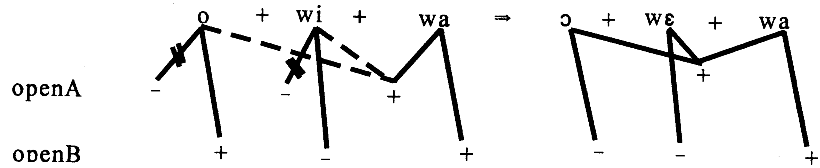

The incongruity of *rúwá* occurs because both Place and Height spread from the root. Space does not allow for a detailed explication of this issue, but this Place-Height interface parallels similar anomalies in Bantu languages (Luganda, Chichewa, Chithumbuka i.a.) and in Turkish. This issue deserves very involved research of its own, and is addressed in Salting (1998:143).

To reiterate, the NS model captures harmony processes in Ogori without the need for cleanup rules or constraints within the height component of the language. Rather, the NS model accurately describes the natural classes as they are delineated in Ogori harmony.

5. Conclusion

For the most part, distinctive features have been assumed to reference fixed phonetic correlates. The Nested Subregister model provides evidence for a linguistic universal in the organizational paradigm rather than in the phonetic referent of universal distinctive features. We have seen the power of the NS model to capture natural classes of segments as they manifest in harmony patterns in two languages, Lhasa and Ogori. Salting (1998) uses the Nested Subregister model to account for harmony in Kɔnni and Somali as well. These are both languages which have proven very problematic to analysis with other models, but which were explained in a straightforward manner when the NS model was applied. Cursory looks at other [ATR] systems such as Yoruba, Akan and Okpe have shown great promise for the NS model as well.

The implication is that the symmetrical, hierarchic division of the vowel height space is a universal, adoptable by some languages. Other languages may employ a gradient division of the vowel height space such as that in Clements (1991), producing scalar harmony patterns. Salting (1998) shows that closely related languages can vary as to which type of organization they choose, as is seen in Bantu languages and dialects of Italian. A prediction is that one would never find scalar harmony in a language which also exhibits evidence for a hierarchical division, and *vice versa*.

OT works in vowel harmony (Pulleyblank 1994) still refer to the traditional features and thus, still require the same sorts of constraints for individual harmony processes. Within the framework of the NS model, there is no gradient violability within the height component. Rather, harmony is seen as a reflection of an internal organization of the inventory. OT may be of benefit in accounting for Height - Place interactions as we saw in Ogori, but this will have to wait on further research.

Perhaps phonological distinctive features (for vowel height at least) are further removed from phonetic correlates than has been assumed, and instead reference a small variety of fixed parameters of organization. The fact that the NS model can so cleanly account for harmony systems which prove very problematic for other systems hints at the need for rethinking our approach to distinctive features. Further research will tell.

REFERENCES

Archangeli, Diana and Douglas Pulleyblank. 1994. Grounded Phonology. MIT Press.

Chang, Kun and Betty Shefts. 1964. A Manual of Spoken Tibetan. University of Washington Press. Seattle.

Chang, Kun and Betty Shefts Chang. 1968. Vowel Harmony in Spoken Lhasa Tibetan. Bulletin of the Institute of History and Philology Academia Sinica. 40:53-124.

Chomsky, Noam & Morris Halle. 1968. The Sound Pattern of English. Harper & Row. New York.

Chumbow, Beban S. 1982. Ogori Vowel Harmony: An Autosegmental Perspective. Linguistic Analysis 10:1, 61-93.

Clements, George N. 1981. The Hierarchical Representation of Tone Features. Harvard Studies in Phonology. Vol. II. IULC.

Clements, George N. 1991. Vowel Height Assimilation in Bantu Languages. Working Papers of Cornell Phonetics Laboratory. 5:37-74.

Dawson, Willa. 1980. Tibetan Phonology. Ph.D. Dissertation. University of Washington.

Hayes, Bruce. 1986. Assimilation as Spreading in Toba Batak. Linguistic Inquiry v.17: 467-499.

Nornang, Willa. 1978. Tibetan Vowel Process: A Case for the Feature, Advanced Tongue Root. D. Malsch, J. Heard and C. Sleat (eds.) Proceedings of the 8th Annual Meeting of the Western Conference on Linguistics.

Odden, David. 1991. Vowel Geometry. Phonology 8:261-289.

Pulleyblank, Douglas. 1986. Tone in Lexical Phonology. D. Reidel. Dordrecht.

Pulleyblank, Douglas. 1988. Underspecification, the Feature Hierarchy and Tiv Vowels. Phonology 5:299-326.

Pulleyblank, Douglas. 1994. Vowel Harmony and Optimality Theory. To appear in Proceedings of Workshop on Phonology. University of Combra, Portugal. September 1993.

Salting, Don. 1995. Vowel Harmony in Lhasa Tibetan: Evidence for [hi] and [ATR]. FLSM VI, 1:39-49.

Salting, Don. 1998. The Nested Subregister Model of Vowel Height. Ph.D. Dissertation. Indiana University.

Yip, Moira. 1980. The Tonal Phonology of Chinese. Ph.D. Dissertation. MIT.