Positional neutralization in the Contrastive Hierarchy: The case of phonological vowel reduction

Christopher Spahr
Generals Paper
University of Toronto

December 18th, 2012

1 Introduction

In this paper[1] I build on the concept of the contrastive hierarchy as developed by Dresher (2009) in order to account for positional neutralization, with a focus on phonological vowel reduction. I propose that one of the main roles that the contrastive hierarchy plays in phonology is in formalizing the patterns in which members of the inventory neutralize. This is done by conceptualizing the contrastive hierarchy as representing the "contrastive structure" of the inventory, such that all nodes of the contrastive tree can be interpreted as phonemes. Segments in reduced positions are represented with non-terminal nodes of the hierarchy, and so instead of neutralization being described as allowing only a subset of the inventory's terminal contrasts, neutralization prevents the use of terminal nodes altogether. I argue that this better captures neutralization as a loss of a contrast, because the feature for which a set of phonemes is contrasted is disallowed rather than being restricted to a single value.

Analysis within this model rests on two core principles regarding phonetic realization. First, phonemic identity (i.e. which node of the hierarchy is represented in a given position) cannot be determined based purely on phonetic realization; the observation that a segment in a position of neutralization sounds like one in a position of full contrast is not sufficient reason to equate a neutralized segment to a non-neutralized one. Second, different phonetic realizations can exist for the same node of the hierarchy; predictable allophony as we know it still exists.

The paper is organized as follows. §2 provides relevant background information on phonological vowel reduction and positional neutralization in general. §3 lays out the non-terminal node model of neutralization that I propose, and applies it to vowel reduction in Bulgarian and Russian. §4 is a case study of European Portuguese which illustrates how predictable allophony can give rise to apparent contrasts which are outside of the scope of positional neutralization. §5 discusses the implications of applying the model I develop in §3 to other kinds of neutralization. §6 concludes.

[1] I would like to thank the members of my generals committee, that is, Elan Dresher, Daniel Currie Hall, and Keren Rice, for all of their advice, feedback, and encouragement to take this work further. It should go without saying that any mistakes and oversights contained herein are entirely my own.

This paper is dedicated to the second mora.

2 Phonological vowel reduction

In this paper, I define phonological vowel reduction as a type of positional neutralization in which the contrast between two or more different vowel phonemes cannot be realized. In other words, segments can be said to have undergone phonological vowel reduction when there is no way to utilize them in contrast with each other in surface forms. This is crucially distinguished from *phonetic* vowel reduction, the differing phonetic realization of vowels in unstressed position, e.g. through centralization, which does not affect the number of contrasts which can be realized in such a position. I consider such non-categorical phonetic reduction to be a kind of allophony, and not a direct concern of the phonology's role in contrast. Nonetheless, both phonological and phonetic vowel reduction can be present in the same language, and the model I argue for below does speak to the relationship between representations in the former and realizations in the latter.

I remain agnostic as to the exact synchronic mechanism compelling phonological vowel reduction, but will assume that some kind of process, either rules in the derivational sense of *The Sound Pattern of English* (Chomsky and Halle 1968) or constraints on surface forms in the sense of Optimality Theory (Prince and Smolensky 2004), results in a smaller number of vowel contrasts to appear in surface representations than can be present in underlying representations. Reduction processes result in a failure to license certain contrasts in a given type of position. Most commonly, for example, vowel reduction is correlated with unstressed positions. For the purposes of discussion here, I will speak in terms of rules changing from one phonemic feature combination to another, applying in the relevant positions, where "phonemic feature combination" refers to those defined in the inventory.

There are some generalizations to be made with regard to the diachronic origins of vowel reduction. Barnes (2006:20) notes that by far the most common kind of contrast neutralized in unstressed positions is that of vowel height. Patterns of positional neutralization of vowels arise due to phonologization of phonetic vowel reduction. Phonetic reduction results from articulatory "undershoot". Due to the shorter duration of unstressed syllables, it is difficult to reach precise articulatory targets for a greater number of contrastive vowels; it is thus considered not to be a coincidence that vowel reduction tends to be a property of languages in which duration is a primary correlate of stress (Barnes 2006:29–30).

According to Barnes (2006:20), the most commonly neutralized contrasts in vowel reduction are height, nasality, and quantity. All three of these have obvious connections to how easily they can be articulated within a durational window. The neutralization of other kinds of contrasts is considered rare or unattested, but in principle a grammar may impose categorical neutralization on any contrast. This is not a problem for Barnes, as he takes an approach similar to Evolutionary Phonology (Blevins 2004), which is not concerned with restricting the set of possible grammars, but instead attributes the apparently restrictive patterns of attested systems instead to phonetic naturalness in language change. It is not a concern of my model, which, as we will see in §5, is intended to deal with other kinds of neutralization in addition to phonological vowel reduction. Restricting the model to features specific to vowel reduction processes would interfere with the more general claims and predictions made about the way neutralizations are phonologically interpreted.

2.1 Some examples of vowel reduction

Phonological vowel reduction is particularly common in Romance and Slavic languages. Barnes (2006:21) gives a "commonly cited example" of height-neutralizing contrast, which comes from Central Eastern Catalan, in which a seven-vowel inventory is reduced to a three-vowel inventory in unstressed syllables.

(1) Central Eastern Catalan vowel reduction (Barnes 2006, adapted)

a.

Stressed	→	Unstressed
i	→	i
e, ɛ, a	→	ə
u, o, ɔ	→	u

b.

r̃íw	'river'	r̃iwɛ́t	'river-DIM'
néw	'snow'	nəwɛ́tə	'snow-DIM'
mɛ́ɫ	'honey'	məlɛ́tə	'honey-DIM
pálə	'shovel'	pəlɛ́tə	'shovel-DIM'
r̃ɔ́ðə	'wheel'	r̃uðétə	'wheel-DIM'
mónə	'monkey-FEM'	munɛ́tə	'monkey-FEM-DIM'
kúrə	'cure'	kurɛ́tə	'cure-DIM'

The vowel correspondences between stressed and unstressed syllables in Central Eastern Catalan are given in (1a). Evidence for these correlations comes from the fact that their phonetic (and contrastive) realizations vary systematically with different stress positions. Examples of the same roots with different stress are given in (1b), where the addition of a diminutive suffix causes the stress to move, and therefore different members of the full set of contrastive vowels to be realized.

Another language which is well known for its phonological vowel reduction is Bulgarian. Stressed syllables in Bulgarian show six contrastive vowels, shown in (2):[2]

(2) Bulgarian stressed vowel inventory

	front	central	back
	non-round		round
high	i		u
mid	e	â	o
low		a	

In unstressed syllables, the six-vowel system can be neutralized in three pairs, depending on the variety of Bulgarian, with phonetic values roughly as shown in (3).

[2]The data here are based on Scatton (1984:54), but with some modifications: The vowel written here as <â> is described by Scatton as high. However, in light of descriptions by Scatton (1975), Crosswhite (2001), Zec (2002), and phonetic data from Tilkov (2009), I regard it here as mid. According to Scatton, <â> is the symbol used in official transliterations of Bulgarian, and so I have adopted it in this paper, following Barnes (2006).

(3) Neutralizing vowel pairs in Bulgarian

â, a	→	a
o, u	→	u
e, i	→	i

The Bulgarian data will be discussed in greater detail below (in §3.1), as will the reduction patterns of Russian (in §3.2) and Portuguese (in §4).

2.2 Distinguishing the phonological and the phonetic

Crosswhite's (1999) dissertation (subsequently published as Crosswhite 2001) develops an oft-cited theoretical approach to vowel reduction. Crosswhite's analysis relies on positing two different kinds of vowel reduction: "contrast-enhancing" and "prominence-reducing". Contrast-enhancing reduction happens in unstressed position in order to realize a smaller number of contrasts in a more dispersed inventory, because it is easier to distinctly produce and perceive smaller sets of contrasts in less prominent positions. The result is that only peripheral vowels can appear in reduced positions. Prominence-reducing reduction, on the other hand, is a tendency to avoid more sonorous vowels in less prominent positions. This follows from the tendency of /a/ to be realized as the less sonorous [ə], rather than the more peripheral [a] in unstressed positions.

The constraints governing the two kinds of reduction can interact in a single grammar in order to give two different degrees of reduction in a language like Russian. In Russian, the five-vowel inventory /i, u, e, o, a/ is reduced to three vowels in unstressed position, /i, u, a/, but this /a/ can be realized in two different ways. "Moderate" reduction occurs in e.g. the (more prominent) pretonic position, where /a/ is realized as the more peripheral [a]. "Radical" reduction, on the other hand, occurs in other positions, where /a/ is realized as the less sonorous [ə].

Barnes (2006) argues against this approach, however. Crosswhite distinguishes between whether an output is the result of contrast-enhancing or prominence-reducing reduction based on the phonetic realization of the unstressed /a/ (which is actually the neutralized counterpart of /a/ and /o/). In order to use this as a diagnostic it would be necessary to draw a (presumably arbitrary) line between these two realizations. But in reality there is no such line to be drawn. Barnes (2006:65) comes to the conclusion that, while merger between /a/ and /o/ in unstressed position is categorical, "[g]radient raising of /a/∼/o/ toward [ə] occurs in all unstressed syllables as a function of the duration allotted to the vowels in question by the phonetics." Longer instances of the vowel will have a lower height, regardless of position; pretonically, it can be closer to [ə] if the vowel is shorter due to speech rate, and in positions of "radical" reduction, it can be closer to [a] when it has a longer duration. Rather than being due to two different kinds of reduction processes, positions of "moderate" reduction tend to have longer duration due to independent prosodic reasons, and the traditional description follows from this.

In other words, Crosswhite has conflated the notion of phonetic (though perceptible) quality with that of phonologically contrastive categories. While the positional neutralization of the /a/–/o/ contrast is certainly categorical, as she describes it, there is no categorical distinction between the two heights at which this neutralized pair can be realized, and thus there cannot be two distinct processes determining them categorically; the difference is merely predictably and gradiently allophonic.

2.3 Inventories and sub-inventories

If positional neutralization such as phonological vowel reduction results in a smaller inventory, one question we must ask is, how is this reduced (sub-)inventory determined? Given that this kind of neutralization is reflected in many-to-one relationships, where several members of the stressed inventory correspond to a single member of the unstressed inventory, it makes sense that a sub-inventory be derived from the full inventory. We must then ask how members of the sub-inventory are represented, and how these representations relate to those of the full inventory.

It should be noted that the present study sets aside cases which resemble neutralization in that there is a smaller number of contrasts which can be realized in certain positions, but where there are no alternations of these positions, such that correspondences between the full inventory and the reduced inventory cannot be seen. In such systems, the more restricted inventory may be used only in certain morphemes, and so can be considered to be completely separate, bearing no direct phonological connection to the larger inventory. Dyck (1995) takes such an approach in her analysis of the vowel inventories of a number of dialects of Spanish and Italian. In many of these varieties, the number of vowel phonemes available in inflectional suffixes (desinences) is smaller than that available in roots. She argues that the differing feature specifications of the members of the desinential inventory must be underlying, as they are not predictable or the result of a phonological process (Dyck 1995:28). The evidence for whether they are different comes from phonological activity: Many Romance varieties have a kind of stem vowel harmony process known as "metaphony", but metaphony can only be triggered by desinential high vowels in desinential vowel systems containing at least four vowels. In systems with fewer than four vowels, it is not necessary to contrastively specify high vowels with the feature [high], and accordingly, they do not act as triggers to high vowel harmony.

Dyck's analysis, while providing insight about the contrastive status of vowels in relation to each other within a sub-inventory, does not apply directly to neutralization; there is no active reduction taking place, as there is no way to see how members of the different inventories correspond to each other, which there is in vowel reduction where the same vowels can alternate between being stressed and unstressed. The separate inventories are truly separate from each other.

For systems in which there are positional alternations that show correspondences between full and reduced inventories, there are two different approaches which can be taken with regards to how these correspondence relations should be captured. The first approach is to assume that the sub-inventory which appears in neutralizing positions is very literally a subset of the full inventory. Crosswhite (2001) assumes this at least implicitly in stating that one phoneme neutralizes to another, where correspondences between the underlying and surface inventory are governed by differently ranked faithfulness constraints, as do authors like Scatton (1984), by saying that Bulgarian /i, u, e, o, â, a/ reduce to /i, u, a/. The problem with such an approach is that the diagnostic for which phonemes are neutralized to is, barring independent evidence from phonological activity, a phonetic one. To say that /e/ and /i/ neutralize to /i/ means that the contrastive phoneme /e/ is disallowed in a position, but /i/ is allowed. The literal subset approach directly equates phonetic realization with phonemic identity. If one of the neutralized vowels sounds like one of the full vowels, then it is the same phonologically. In some sense, however, neither of these phonemes is allowed in neutralized position, as neither remains fully in contrast with the other. The phonology of neutralization can better capture this if it does not use members of the full set of contrasts in neutralized position.

We saw above that as Barnes (2006) points out, the literal subset approach can lead to trouble when it comes to defining categorical reduction. Crosswhite correctly recognizes a categorical neutralization of /a/ and /o/ in unstressed position, but mistakenly identifies reduction to a vowel which is phonetically different from the members of the fully contrastive set ([ə]) as having a different phonological status than the other unstressed realization of this pair ([a]). In other words, a sound is considered to be phonologically different simply because it is perceptibly different, despite the fact that (as Barnes shows) it is gradiently and predictably realized allophonically, and there are no categories to speak of. I would argue, thus, that it is important to keep in mind the difference between contrast and differing phonetic realization when analyzing reduction in vowel inventories.

The other approach to neutralization is to say that the phonemic status of a neutralized pair is archiphonemic. This is related to the Prague School concept of the “archiphoneme” (see Davidsen-Nielsen 1978 for extensive discussion of this concept). The idea of the archiphonemic approach is that positions corresponding to the neutralization of a given pair of the full contrastive set is somehow a combination or common denominator of that pair, or a lack of specification for the feature which minimally distinguishes that pair. I adopt this sort of approach in the present work. I lay out my model of neutralization, as it relates to Dresher’s (2009) Contrastive Hierarchies, in the next section.

3 Inventories of neutralization

Trubetzkoy (1969:228) said of neutralizations that "They are just as characteristic of the phonemic system of the individual languages and dialects as are the differences in the phonemic inventory." Because neutralization is by definition a loss of contrast, I hold patterns of neutralization to be one of the core concerns of defining the contrastive structure of the inventory. I will assume the notion of the Contrastive Hierarchy (Dresher 2009) as a conceptualization of the contrastive structure of the inventory. In this model, language-specific feature orderings are applied to inventories, yielding branching tree diagrams that show the contrastively specified features for a given phoneme.

Let us take, for example, a language with a simple three-vowel inventory. A contrastive hierarchy can reflect phonological activity depending on the features it assigns and the order it assigns them in. Consider the inventory in (4):

(4)

	FRONT	BACK/ROUND
HIGH	/i/	/u/
LOW	/a/	

Suppose that /u/ clearly patterns as a contrastively round vowel, perhaps causing phonological alternation by spreading a [+round] feature. This can be reflected in a binary branching tree which orders the feature [±round] first, dividing off /u/ as [+round]. The remaining segments, which are not contrastively [+round], are divided in another binary split under the [–round] branch. If for example, /a/ is patterning as a low vowel and /i/ as a non-low vowel, then this divide can be between [±low]. This is shown in (5a).

(5) Three-vowel hierarchies with [±round] ordered first

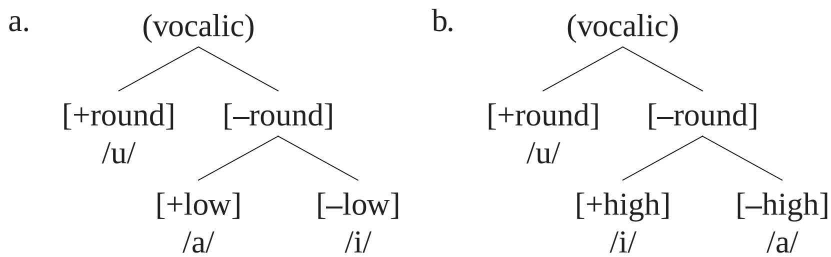

If on the other hand /i/ is patterning as a high vowel to the exclusion of both /a/ and /u/, then the second contrastive feature can instead be [±high], as in (5b). The power of this system lies in the ability to change the ordering of features in order to reflect different classes based on groupings of phonological features. For example, if /u/ and /i/ both group together as high vowels to the exclusion of /a/, then the feature [±high] can be understood to be ordered first in the hierarchy, as in (6). Here /a/ does not receive a contrastive value for any feature other than [–high], because it has no other feature to contrast with. Thus the use of contrastive hierarchies predicts that in a system such as that in (6), /a/ will not be able to act as phonologically [+low], and so this means of assigning representations limits the number of contrastive features that can be used depending on the number of contrasts in a system.

(6) Three-vowel hierarchy with [±high] ordered first

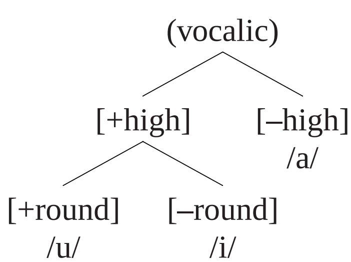

In the model that I propose in this paper, the ways that the terminal nodes of a contrastive hierarchy group together is key to reflecting patterns of neutralization. In a weak interpretation of Dresher's (2009) use of the contrastive hierarchy, the tree diagrams serve merely as a notation for the order in which an inventory has been algorithmically divided into contrastive features. I propose here that the hierarchy itself is in a sense a real phonological object, and that it represents what I call the internal "contrastive structure" of the inventory. We can refer to this as taking a "strong interpretation" of the contrastive hierarchy. One of its key roles is in determining the patterns and phonological representations used under neutralization. My main claim is that neutralized contrasts are represented with non-terminal nodes of the hierarchy, and so sisterhood relations within the tree, not only the features assigned, matter.[3] In this way, all nodes of the hierarchy are viable as members of the inventory, and are thus phonemes (or much like the Prague School sense of "archiphonemes").

There are several major theoretical advantages to this model. It provides a restrictive and principled set of possible neutralizations; once a hierarchy is established, pairs can only neutralize to non-terminal nodes by which they are immediately dominated.[4] To put it another way, the patterns of neutralizations in the phonology give us clues as to the shape of the contrastive hierarchy. The featural makeup of these non-terminal nodes then makes predictions about the possible phonetic realizations of neutralized pairs. For example, if the phonemes /u/ and /i/ from the hierarchy in (6) neutralize, and our hierarchy shows that this involves neutralization for the feature [$\pm$round], then the non-terminal node dominating /u/ and /i/ will have no specification for [$\pm$round]. The realization of this node will then not necessarily have to be within phonetic space defined by a contrastively [+round] vowel (like /u/). We would then predict that the neutralization of these vowels might be closer to [ɨ], all things being equal.

The other advantage of this model is that the non-terminal nodes of reduction can be used in underlying representations for morphemes which show no stress alternations. To take the /u/–/i/ example from above, if we have a vowel which never surfaces as stressed, and is instead always heard by a learner as [ɨ], then there is no need to arbitrarily posit an underlying representation containing /u/ or /i/, either of which could have been the underlying source. Instead, the underlying

[3]Several recent papers (Harvey 2012, Oxford 2012) have pointed out the importance of sisterhood relations within the hierarchy in predicting phonological mergers in diachronic change, but do not necessarily motivate their importance in synchronic grammars.

[4]It can in fact be said more generally that a set of terminal nodes can only be neutralized to a non-terminal node that exhaustively dominates its members.

representation can be made more economical and less complex by simply specifying the non-terminal node, in absence of any evidence pointing to either of the terminal nodes.

When analyzing data within the model I am proposing, there are two main principles which must be kept in mind with regards to the relationship between phonetic realization and phonemic representation: (i) just because two segments sound the same, it does not mean that they are the same phonologically, and (ii) just because two segments sound different, it does not mean that they are different phonologically.

(i) means that neutralization of /u/ and /i/ in this example should not be understood as all instances of /u/ changing to /i/ in neutralizing position simply because the phonetic realization [ɨ] sounds relatively close to a stressed /i/ realized as [i]. [ɨ] is the phonetic realization of the non-terminal node dominating /u/ and /i/; it need not be the realization of an unstressed /i/ specifically, so much as the realization of the node /u/–/i/.

(ii) refers simply to allophony. When allophones are perceptually different to a person doing a phonological description, there is a tendency to consider the difference in realizations to be relevant to phonemic status, despite the fact that there is no loss of contrast. As we will see with Russian below, my model allows for non-terminal nodes of the hierarchy, which have their own phonemic status, to show allophony independently from terminal nodes. This does not prevent allophones of terminal and non-terminal nodes from sounding similar or the same, but phonetic realization is not necessarily relevant to phonemic identity.

Next I will show how the principles of the model I have sketched here apply to the positional neutralization of vowels in Bulgarian and Russian.

3.1 Bulgarian

Recall that stressed syllables in Bulgarian show a six-vowel inventory, as shown in (7) (repeated from 2 above):

(7) Bulgarian stressed vowel inventory

	front	central	back
	non-round		round
high	i		u
mid	e	â	o
low		a	

As mentioned above, the central, back, and front vowels neutralize in three pairs depending on dialect: /â/–/a/, /u/–/o/, and /i/–/e/. Scatton (1984:57) notes however that these pairs neutralize in a "rigid hierarchy", depending on variety and register. Virtually all dialects and registers will neutralize unstressed /â/ and /a/, which is realized as [ə]. Some dialects will further neutralize /u/ and /o/, whose merger is realized as [ʊ]. Finally, some "non-literary" varieties will neutralize unstressed /i/ and /e/, the result of which is realized as [ɪ]. This is schematized as the implicational hierarchy in (8); any variety which neutralizes /i/–/e/ must also neutralize /u/–/o/ and /â/–/a/, but a variety which neutralizes /u/–/o/ does not necessarily also neutralize /i/–/e/.

(8) Implicational hierarchy of Bulgarian vowel pair neutralizations

/i/–/e/ > /u/–/o/ > /â/–/a/

This means that unstressed syllables in Bulgarian dialects show three possible contrastive vowel inventories, which are given in (9):

(9) Possible unstressed inventories in Bulgarian

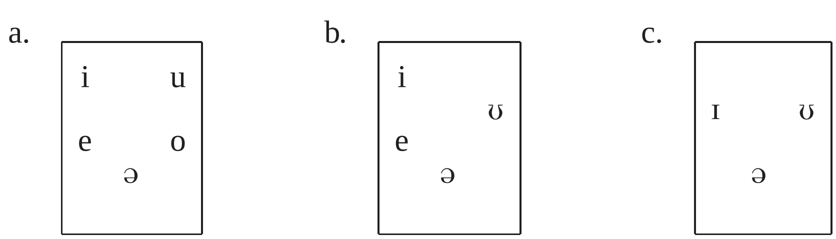

Let us now turn to the way the inventories are reflected in a contrastive hierarchy. Considering first the full inventory of stressed syllables in (7), we know that we will need to distinguish three main classes of vowels: front vowels, back/round vowels, and central vowels. The patterns of vowel reduction clearly show that these three pairs form classes, but there is independent evidence in Bulgarian phonology which points towards such a classification. For example, the front vowels, to the exclusion of non-front vowels, disallow a contrast between palatalized and non-palatalized consonants proceeding them (Scatton 1984:64–65). Let us consider [±front] to be the feature at the top of the vowel hierarchy.[5] This has the effect of separating off /i/ and /e/ from the rest of the inventory.

The pairs /u, o/ and /â, a/ can then be distinguished from each other by the feature [±round]. The height distinctions within the three pairs can be distinguished with the feature [±high] for the front and round vowels, and [±low] for the mid vowels. We thus arrive at the tree in (10):[6]

(10) Contrastive hierarchy for Bulgarian

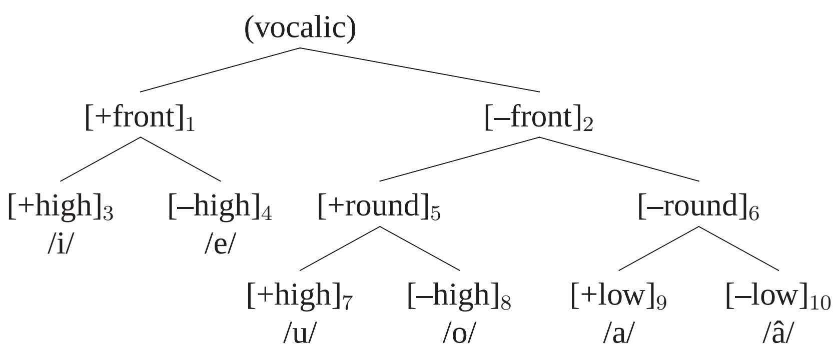

The terminal nodes in this tree account for the contrastive distinctions between the six stressed vowels of Bulgarian, but how are the sub-inventories in (9) represented? When the pair /â/–/a/ is neutralized, unstressed positions can no longer show a contrast between nodes 9 and 10 of the hierarchy. In other words, [–round] segments in unstressed position are prohibited from having a contrastive value for the feature [±low]. Instead of both nodes neutralizing to node 9 or to node 10, a rule changes both nodes 9 and 10 into node 6. Node 6, as a non-terminal node, has its own phonemic status as a position in the structure of the inventory in my model. It has its own feature specifications, [–front, –round], and its own phonetic realization: [ə]. This realization

[5] I will use binary rather than privative features in the feature hierarchies in this paper, setting aside for the time being the debate between binarity and privativity.

[6] I have labelled the nodes of this tree with subscript numbers for the purposes of discussion.

follows from the fact that it has no height specification; being specified as a non-front heightless vowel, it is realized as mid-central. Thus the differing phonetic realization follows directly from implementation of the feature specifications predicted by a contrastive hierarchical model.

In some varieties, there is also a neutralization of the round (back) vowels in unstressed position. Among the [+round] terminal contrasts, that is, nodes 7 and 8, a contrastive value for [±high] cannot be realized, and so the pair /u/–/o/ is merged. Again, here, a rule changes both nodes 7 and 8 to node 5. Node 5 also has its own phonemic status and feature specifications, [–front, +round], and its own realization, [ʊ]. As with the realization of the neutralized /â/–/a/ pair, the lack of a specified height feature causes the (positionally only) round vowel to be realized lower than if it were contrastively [+high].

The last neutralized pair, /i/–/e/, is accounted for in an analogous way. In unstressed positions, nodes 3 and 4 are both changed by rules to node 1, causing them to lose their contrastive value for [±high]. Node 1 has the feature specifications [+front], and is realized as [ɪ]. Once again, by having no height specification, node 1 is realized lower than it would be if it were specified.

The tree in (10) also captures the implicational hierarchy in (8). If we assume that positive values for binary features represent the marked member of an opposition (either in a universal or language-specific way), then we notice a pattern in the three neutralizing sub-trees in (10). The first pair to neutralize, /â/–/a/, is dominated by two unmarked feature values, [–front, –round]. The next (/u/–/o/) is dominated by one unmarked and one marked feature value, [–front, +round]. Finally, /i/–/e/ is dominated only by a marked feature value, [+front]. While the model is not necessarily intended to reflect implicational hierarchies such as this, it is a nice result of the analysis that it is reflected. Whether this is a more general tendency in contrastive hierarchies of neutralization will be a subject of future research.

3.2 Russian

Stressed syllables in Russian show a five-way vowel contrast:[7]

[7] I follow Iosad (2012:523) in assuming that the mid-high vowel [ɨ] is not a phoneme, but rather is in complementary distribution with [i] as an allophone of /i/ depending on palatal context. For some discussion of this issue, see for

(11) Russian stressed vowel inventory

i		u
e		o
	a	

As discussed briefly above in §2.2, traditional descriptions of Russian vowel reduction show two degrees of vowel reduction. These refer to the way that the merger of the pair /a/–/o/ is realized phonetically in unstressed positions. While there seems to be no disagreement over the fact that in varieties which have phonological vowel reduction, the neutralization of /a/ and /o/ is complete, there is disagreement over how to treat the surface realization of these segments in unstressed syllables. The two are said to be realized together as [ə] in positions of "radical" reduction, but positions of "moderate" reduction result in a lower realization, closer to [ɐ] or even [a]. Iosad (2012:524) gives the following positions as those of moderate reduction:[8]

(12)
- The syllable immediately preceding the stressed one (the "first pretonic");
- Onsetless syllables, regardless of stress (though this is somewhat contested);
- Gradient effects in phrase-final unstressed open syllables;
- Some claim both vowels in a hiatus will undergo moderate reduction.

Crosswhite (2001) chooses to analyze a distinction between moderate and radical reduction very literally, motivating them with two different kinds of constraints for two different kinds of reduction. Raising of the low vowel occurs due to a preference to avoid sonorous vowels in less prominent positions called "prominence-reducing" reduction. However, doing so presupposes that the realizations [ə] and [ɐ] of the unstressed /a/–/o/ pair fall into categories. Barnes (2006) shows that they do not form categories, but that instead the realized height of the pair in unstressed posi-

example Timberlake (2004, §2.2.2).

[8] I have changed the wording somewhat from the original source for reasons of conciseness.

tion is a very predictably gradient matter, varying as a function of duration regardless of position within a word.

While neutralization of /a/–/o/ in all unstressed contexts is categorical, Iosad (2012) notes that the realization of unstressed vowels does depend on several variables: whether they are in a context as described in (12) above, and whether they are in a "palatalized context", that is, following a palatalized consonant. Non-palatal moderate reduction results in a reduction from a five-vowel inventory (of stressed positions) to a three-vowel inventory, with the phonetic realizations in (13), via the contrastive merger of the pairs /i/–/e/ and (as discussed above) /a/–/o/. The phoneme /u/, while centralized phonetically, remains contrastively distinct from the rest of the vowels.

(13) Moderate reduction in a non-palatal context (Iosad 2012:526)

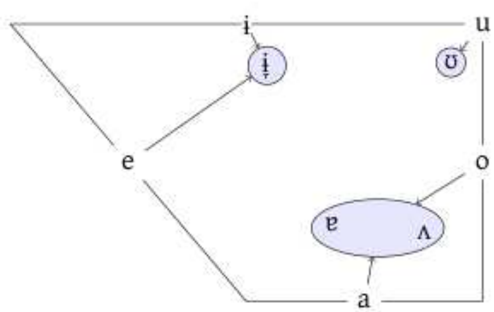

All other contexts result in reduction to a two-vowel system, where /u/ remains distinct (as in 13), and the other four vowels, /i, e, o, a/ are realized the same phonetically. In palatalized moderate context, /u/ again centralizes slightly to [ʊ] but remains distinct, while the other four are realized as a fairly high-front vowel [ɪ]:

(14) Moderate reduction in a palatal context (Iosad 2012:527)

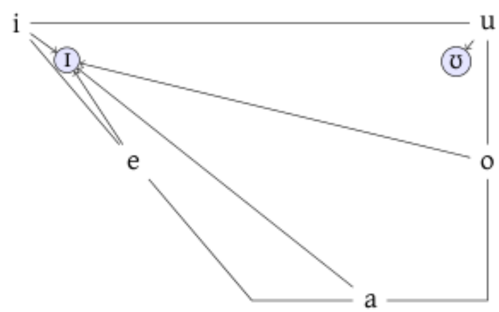

In both palatal and non-palatal radical reduction contexts, /u/ is the only phoneme which retains its contrast. As seen in (15), the only difference is the realization of neutralized /i/–/e/–/o/–/a/, which is more high and front in a palatalized context.[9]

(15) Radical reduction (Iosad 2012:529)

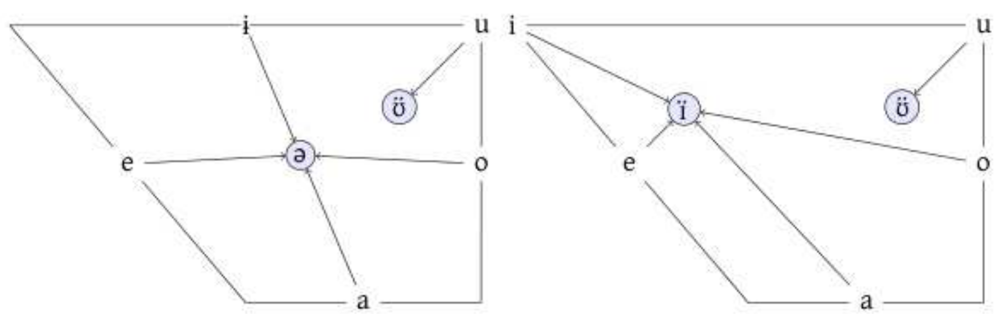

(a) Non-palatalized context (b) Palatalized context

The phonological (and phonetic) reduction data show several classes which must be captured by a contrastive hierarchy. The most obvious fact about Russian vowel reduction is that, regardless of the kind of context, /u/ remains distinct; it tends not to undergo merger with any other phoneme. Iosad (2012:532) notes however that "in casual or fast speech it can lose its labialization and neutralize with other vowels in a [ə]- or [ɪ̈]-like sound". Even when it does not lose labialization, its realization "varies very widely, from [u] to [ʊ] and all the way through [ʊ̈] to [$ə^w$]." Nonetheless, considering it is both the last thing to undergo merger, and it retains in its weakest distinct realiza-

[9]Iosad uses a diaeresis (¨) to distinguish whether a phonetically reduced vowel occurs in a moderate or radical reduction context.

tion ([əʷ]) some minor amount of labiality, it makes sense to assume that its contrastive feature is [+round], and that this is the first division to be made in the tree.[10]

/i/ and /e/, being phonetically front vowels in their non-neutralized realizations are both then divided off by the feature [+front], where they form a class which merges contrastively in moderate non-palatal reduction contexts. In non-reducing contexts, they are distinguished from each other by the feature [±high]. The remaining two vowels, /o/ and /a/, are grouped together as [–front], and themselves form a class of categorical merger in unstressed contexts. They are distinguished from each other by the feature [±low]. We thus arrive at the contrastive hierarchy in (16):

(16) Contrastive hierarchy for Russian

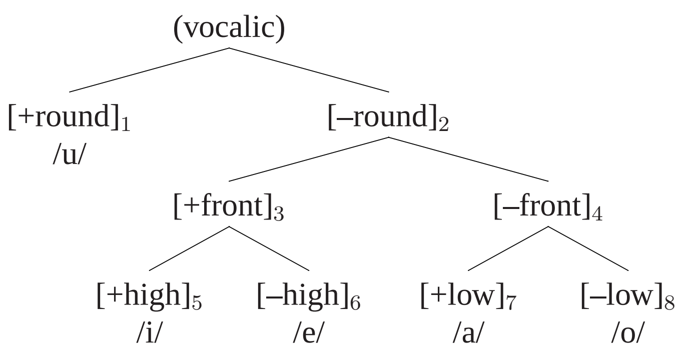

It will be apparent that the tree in (16) specifies the terminal node labelled /o/ as contrastively [–round]. Iosad (2012:537–538) holds that /u/ and /o/, which are both phonetically labial, in no way pattern together. In fact, according to him, /o/ does not pattern as phonologically labial in any processes. Because /u/ is clearly distinctively round, and since it maintains its own contrast to the exclusion of the rest of the inventory in all reduction contexts, it is a requirement of the model I put forth here that /o/ not be specified as [+round]. It is thus striking to note that it in turn predicts exactly the kind of phonological specifications that Iosad ends up using, including this idiosyncratic property of /o/. Iosad works within the Parallel Structures Model (PSM), which only specifies features that are phonologically active. Using privative features, Iosad specifies the

[10] Recall from the discussion in §3.1 that the implicational hierarchy of Bulgarian neutralizations can be reflected by placing the last contrastive mergers higher in the contrastive hierarchy. This approach applies here, as well, as /u/ neutralizes only after all of the other vowels have.

Russian vowel inventory as in (17):

(17) Russian privative feature specifications in the PSM (Iosad 2012:538)

	V-manner		V-place	
Vowel	[open]	[closed]	[labial]	[coronal]
/a/	✓			
/o/		✓		
/e/		✓		✓
/i/				✓
/u/			✓	

The result is remarkably similar what is achieved using the more restrictive means of feature specification offered by the contrastive hierarchical approach, as seen in (18). The only difference lies in the way that height is treated for the mid vowels /o/ and /e/; Iosad specifies them both as [closed], while my analysis calls for no positive height features for mid vowels at all:

(18) Russian feature specifications in my analysis

Vowel	[±low]	[±high]	[±round]	[±front]
/a/	+		(–)	(–)
/o/	(–)		(–)	(–)
/e/		(–)	(–)	+
/i/		+	(–)	+
/u/			+	

Iosad's use of privative [closed] features on the mid vowels allows him to formalize non-palatal reduction as a deletion of this [closed] feature. In doing so, /e/, which is specified [closed, coronal], becomes simply [coronal], which gives it identical features to /i/, effectively turning it into /i/. In other words, Iosad assumes a literal subset approach to vowel reduction: the correspondences of positional neutralization involve changing from one member of the full inventory to another. However, since I assume the non-terminal node (archiphonemic) view of neutralization, such specifications need not be stipulated. /o/ and /e/ are realized as mid vowels phonetically because they bear no positive specifications for height features at all.[11]

Let us now consider the phonetic implementation of the non-terminal nodes which are used in

[11] The use of [closed] on /o/ causes further complications. When [closed] is deleted in reduction, /o/ loses its only feature, leaving a bare root node. Iosad must assume that bare root nodes are not licit representations in Russian.

unstressed position. Essentially, there can be said to be two different phonological sub-inventories in Russian. In a moderate non-palatal context (cf. 13), there are three different contrastive vowels. /u/ does not neutralize to any other phonemes, and so its representation remains the same, as node 1 of the tree in (16), with the feature specification [+round]. Other vowels in moderate non-palatal contexts lose their specifications for height features, and so rules force a change to nodes without height specifications; /i/ and /e/ (nodes 5 and 6) change to node 3, and /a/ and /o/ (nodes 7 and 8) change to node 4. This gives us the sub-inventory and specifications in (19):

(19) Russian feature specifications in a moderate non-palatal reduction context

Phoneme	Node	Feature Specifications	Phonetic Realization
/u/	1	[+round]	[ʊ]
/i/–/e/	3	[–round, +front]	[ɨ̞̟]
/a/–/o/	4	[–round, –front]	[ɐ, ʌ]

Recall the phonetic realizations in (13), which correspond to these reductions. If we consider the feature specifications in (19) to constitute their own inventory for the purposes of phonetic implementation, then I would suggest, following Hall's (2011) model of emergent phonetic dispersedness, that these contrastive feature specifications lead to exactly the phonetic realization we find. /u/'s contrastive labiality is enhanced with backness, which is most salient at the top of the vowel space, yielding its high-back realization. The pair /i/–/e/ is specified as front, which prevents it from moving too far back. It can move backwards slightly, towards phonetic [ɨ], however, provided it does not move into contrastively [–front] territory, which it does not. Finally, /a/–/o/ is specified as neither front nor round, and lacking height specifications as it does, it is left to occupy a range in the approximate center of the vowel space.

The other sub-inventory derived in Russian vowel reduction contains only two contrastive vowels. This is the inventory which occurs in the moderate palatal context, and in both the palatal and non-palatal radical context. In this two-vowel inventory, underlying /u/ again corresponds to a lack of surface neutralization, and so node 1 of the tree in (16) remains distinct. The other four vowels, /i/–/e/–/a/–/o/, do neutralize, however. Thus nodes 5, 6, 7, and 8 are all changed to node 2 by a rule which disallows contrastive specifications for both height and place ([±front]) features. The sub-inventory thus has the feature specifications in (20):

(20) Russian feature specifications in a moderate palatal or radical reduction context

Phoneme	Node	Feature Specifications	Phonetic Realization
/u/	1	[+round]	[ʊ]∼[ʊ̈]
/i/–/e/–/a/–/o/	2	[–round]	[ɪ]∼[ə]∼[ɪ̈]

The phonetic implementation of this two-vowel inventory's specifications does vary slightly depending on context, however. In an all-things-being-equal context, such as we find with non-palatal radical reduction (cf. 15a), the realizations are exactly as we would expect. /i/–/e/–/a/–/o/ (node 2), lacking any place or height specifications, is realized as the most mid-central vowel possible: [ə]. /u/ is still specified for labiality, which is again enhanced with backness as predicted by Hall (2011), but due to the lesser duration typically afforded in such contexts, it undergoes more significant centralization. Given the fact that it remains labial, this does not interfere with its ability to be contrastive.

Palatalized contexts, as in (14) and (15b), show a high-front realization of /i/–/e/–/a/–/o/, which is not to be expected if they lack place and height specifications. However, because these are in positions following palatalized consonants, this is likely due to coarticulation. In other words, the consonant's feature for palatalization (either [+front] or [coronal]) is spread to the vowel, causing a phonetically fronted allophone of node 2. /u/ on the other hand is not affected as strongly by this fronting because it has a specification for labiality, which is enhanced by backness, counteracting the fronting effect. In fact, allophonic fronting seems to be a characteristic of all vowels in the radical palatal context. According to Alexei Kochetov (p.c.), the /u/ which Iosad transcribes as [ʊ̈] in (15b) is actually somewhat fronted as well. As with fronting of the /i/–/e/–/a/–/o/ node, however, this fronting can be attributed to coarticulation with the preceding palatalized consonant, not to the contrastive features of the vowel.

In this section, I have argued that, in cases where two vowels undergo phonological neutralization, neutralization does not take place from one member of the fully contrastive inventory to another, but rather to a third member, represented in the structure of the inventory as non-terminal nodes of the contrastive hierarchy. Thus when, in Bulgarian, /u/ and /o/ neutralize to be realized as [ʊ], it is not that all instances of /o/ become /u/, and /u/ is allophonically centralized in unstressed position to be realized as [ʊ], but rather that both /u/ and /o/ are changed to the non-terminal node. Because this non-terminal node is not contrastively specified for a feature [±high], its centralized realization (as well as its property of being the result of neutralization) follows from its featural representation. Hence, the fact that the unstressed realization of /u/–/o/ is similar to that of stressed /u/ does not require that they are both phonologically /u/. The contrastive hierarchy provides non-terminal nodes, which serve as additional members of the inventory which can be used to represent the archiphonemic segments of reduction, even underlyingly, as with non-alternating reduced vowels.

We have also seen that non-terminal nodes, as phonemes, can have their own phonologically conditioned allophony, as was the case with the palatal versus non-palatal realization of the non-terminal node corresponding to Russian /i/–/e/–/a/–/o/. In the next section I will provide a case study of phonologically conditioned allophony of a terminal node showing an apparent surface contrast, as is the case with European Portuguese /a/.

4 Surface-contrastive allophony: The case of Portuguese

In Portuguese, the vowel /a/ does not undergo neutralization in any position, but it does undergo raising to [ɐ] in unstressed syllables, and in stressed position before nasals and palatals. In this way it is similar to the Russian /u/, which is apart from the rest of the neutralizations in vowel reduction, undergoing only allophonic centralization. Furthermore, the realizations of Portuguese /a/ varying in height between [a] and [ɐ] appear to be a function of duration, much like the realization of the neutralized pair /a/–/o/ in Russian. The difference is that, despite being allophones of a single phoneme, [a] and [ɐ] show a surface contrast in European Portuguese /a/-theme verbs, yielding pairs like [fɐlɐ́muʃ] *falamos* ‘we speak’ [fɐlámuʃ] *falámos* ‘we spoke’.

I will propose that /a/ raising is blocked when the duration which would cause a lower realization is contributed by the presence of two V timing slots on the vowel. The apparently categorical effect in verb paradigms is due to the presence of a floating V timing slot which serves as the past tense morpheme. The use of two timing slots is evidenced by allophonic blocking of raising upon the coalescence of two unstressed instance of /a/ known as “crasis”. Nowhere do the underlying features of /a/ come in to play, and there is no need to posit a separate /ɐ/ phoneme or paradigm-specific allophonic rules. Most importantly, this analysis makes several testable predictions about the relative phonetic durations of the segments involved.

I begin by describing the vowel inventory and patterns of vowel reduction in Brazilian and European Portuguese, followed by working out the contrastive hierarchy and terminal-to-non-terminal node correspondences which account for the reduction data within the model which I outlined in the previous section. I then give my analysis of the blocking of /a/ raising in which the presence of two V timing slots prevents raising.

4.1 Vowel reduction

Most sources agree that both Brazilian and European varieties of Portuguese contrast seven oral vowels in stressed syllables, as shown in (21):

(21) Portuguese stressed oral vowels

	front		back
high	i		u
upper mid	e		o
lower mid	ɛ		ɔ
low		a	

According to Wetzels (2010), in Brazilian Portuguese unstressed syllables, the contrast between the upper mid and lower mid vowels is neutralized, so /e/–/ɛ/ and /o/–/ɔ/ merge contrastively, while the other vowels remain distinct. In word-final open syllables and those closed by /s/, the contrast between all of the non-low front vowels (/i/–/e/–/ɛ/) and non-low back vowels (/u/–/o/–/ɔ/) is neutralized, with high realizations. This is summarized in (22), with the phonetic values of the reduced vowels adapted from Barbosa and Albano (2004).

(22) a. Brazilian Portuguese unstressed syllables

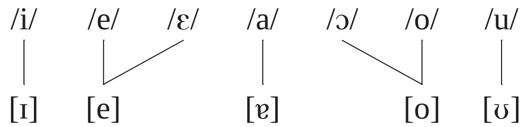

b. Brazilian Portuguese unstressed word-final syllables

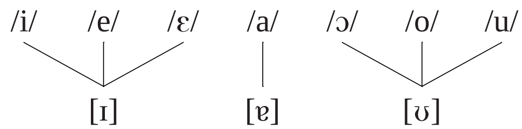

Carvalho (2011:7) points out that the situation is slightly different in European Portuguese, in that unstressed back (round) vowels neutralize from three degrees of height to one, while the mid front vowels /e/–/ɛ/ neutralize to some kind of mid vowel, leaving the high-front vowel /i/ without undergoing any neutralization. This is summarized in (23).[12]

(23) European Portuguese unstressed syllables (Carvalho 2011)

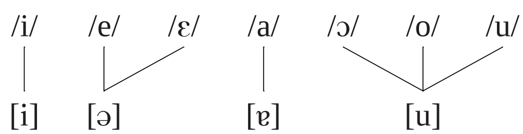

According to Mateus and d'Andrade (2000:18), the unstressed /e/–/ɛ/ pair in European Portuguese is actually realized as a high central vowel [ɨ], and is frequently deleted in colloquial speech, sometimes yielding clusters, so that *dever* [dɨvéɾ] 'duty' can be realized as [dvéɾ], and bate [bátɨ] 's/he beats' can be realized as [bát]. They note furthermore that word-final unstressed [u] can also be dropped, such that *bato* [bátu] 'I beat' can be pronounced [bát] as well, but this is not as common.[13]

[12] I have changed Carvalho's (2011) transcription of the unstressed realization of /a/ from [ɜ] to [ɐ] throughout this paper to match that of the other sources.

[13] It is unclear whether this applies to all of /u/–/o/–/ɔ/ or to only a subset of the three underlying phonemes which can surface as [u] in unstressed position.

4.2 A hierarchy

The first thing to notice about European Portuguese vowel reduction is that, in all contexts and all varieties, /a/, despite being phonetically raised to [ɐ] in unstressed positions, fails to undergo any neutralization with other vowel phonemes. In this way it is much like Russian /u/, as discussed in §3.2. With the same rationale, we can say that /a/'s contrastive feature, [+low], is ordered first in the hierarchy, effectively dividing it off from the rest of the inventory.

The remaining vowels, which are all non-low, group together based on roundness, and the fact that they result in contrastive mergers phonologically. The three heights within each of these groups are then divided with the high vowels /u/ and /i/ being specified [+high]. This is because /i/ retains its contrastive value separately from the two front mid vowels /e/ and /ɛ/, which merge together leaving /i/ separately contrastive in European Portuguese, and in non-final contexts in Brazilian Portuguese. The upper and lower mid vowels, specified [–high], can then distinguished by a feature such as [±tense]. From this we arrive at the hierarchy in (24):

(24) Contrastive hierarchy for Portuguese

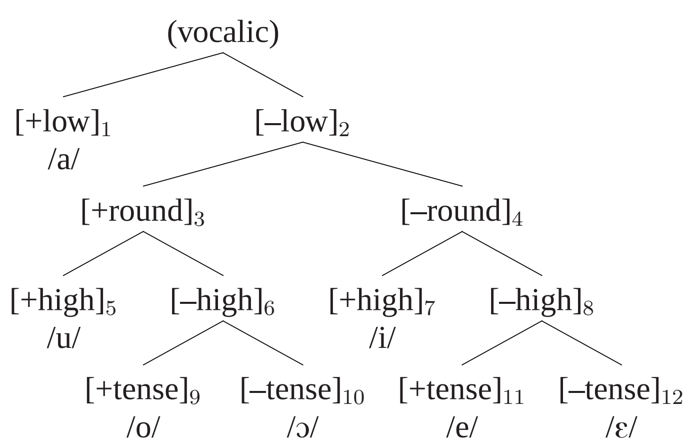

The usual principles of processes governing positional neutralization in my model apply to the Portuguese data here as I demonstrated above for Bulgarian and Russian. In Brazilian Portuguese non-final unstressed syllables, where the /o/–/ɔ/ and /e/–/ɛ/ contrasts are neutralized, contrastive values for the feature [±tense] are not allowed, and so rules change nodes 9 and 10 to node 6, and change nodes 11 and 12 to node 8. Nodes 6 and 8, lacking a specification for [+high], have a fairly mid realization, maintaining their contrastive values for [±round]. In word-final unstressed syllables, where the mid vowels further neutralize with the high vowels, specifications for both [±tense] and [±high] are disallowed, so rules change nodes 5, 9, and 10 to node 3, and likewise change nodes 7, 11, and 12 to node 4.

As can be seen in (25), the five-vowel system of non-final unstressed syllables in Brazilian Portuguese contains specifications for three heights; low ([+low]), high ([+high]), and mid ([–low, –high]).

(25) Brazilian Portuguese feature specifications in non-final unstressed syllables

Phoneme	Node	Feature Specifications	Phonetic Realization
/a/	1	[+low]	[ɐ]
/u/	5	[–low, +round, +high]	[ʊ]
/o/–/ɔ/	6	[–low, +round, –high]	[o]
/i/	7	[–low, –round, +high]	[ɪ]
/e/–/ɛ/	8	[–low, –round, –high]	[e]

The three-vowel system of final unstressed syllables in (26), however, only specifies two heights, with no specifications for [+high]. The non-low vowels are still realized as high, as predicted by Hall (2011), due to enhancement of specifications for [±round]; the backness/frontness which enhances labiality/illabiality is more salient at the top of the vowel space.

(26) Brazilian Portuguese feature specifications in final /V(s)/ unstressed syllables

Phoneme	Node	Feature Specifications	Phonetic Realization
/a/	1	[+low]	[ɐ]
/u/–/o/–/ɔ/	3	[–low, +round]	[ʊ]
/i/–/e/–/ɛ/	4	[–low, –round]	[ɪ]

There are some small differences in European Portuguese, primarily in that the front vowels at no point undergo a three-way neutralization, while the round vowels do. This, as with the implicational hierarchy of Bulgarian neutralizations mentioned in §3.1, is captured by the idea that neutralizations tend to take place first under the (unmarked) negative values of features within the hierarchy. /e/ and /ɛ/ are, in the hierarchy in (24), dominated only by negative feature specifications, unlike any of the other nodes. As for why /e/–/ɛ/ has a tendency to delete in unstressed position rather than merely reducing, the negative values effect may also be responsible: note that /e/–/ɛ/ is the only phoneme in the reduced inventory which has all negative feature specifications (cf. 27), and so we might expect it to tend toward deletion, given its low markedness. In a privative system, it would even be completely featureless, meaning it would be equivalent to a bare root node in unstressed position, a configuration which would be hard to distinguish from the complete lack of a segment in word-final position.

(27) European Portuguese feature specifications in unstressed syllables

Phoneme	Node	Feature Specifications	Phonetic Realization
/a/	1	[+low]	[ɐ]
/e/–/ɛ/	8	[–low, –round, –high]	[ə, ɨ]
/i/	7	[–low, –round, +high]	[i]
/u/–/o/–/ɔ/	3	[–low, +round]	[u]

There are a number of questions which I have not attempted to answer here with regards to the representation of vowel height. Both Wetzels (2010, 1995) and Carvalho (2011) analyze Portuguese as having a four-height system, without the use of features such as [high] and [low]. Wetzels uses inherently hierarchical multiple [open$_n$] features based on Clements (1991), while Carvalho represents height using privative A "resonance elements". I leave investigation of the possibility (and utility) of integrating multiple height features into a contrastive hierarchical model of reduction such as I put forward in the present study to future research. What is important here is the "shape" of the inventory's hierarchical structure in (24), such that there are non-terminal nodes corresponding correctly to the neutralizations which occur in Portuguese phonology.

4.3 "Crasis"

While /a/ does not undergo any neutralization in European Portuguese in any position, as noted above, there are several positions in which it is allophonically raised to [ɐ], including in unstressed position. However, in certain contexts in European Portuguese /a/ raising can be blocked. One example of this is a phenomenon known as "crasis" (Portuguese *crase*), where the coming together of two unstressed /a/s across a word boundary, which would individually be pronounced as [ɐ], form a single vowel pronounced [a]:

(28) Crasis (Carvalho 2011:6, adapted)

a.	VR:	/a/ → [ɐ] in unstressed syllables		
b.	Crasis:	[ɐ] + [ɐˈmigɐ] → [aˈmigɐ]	*a amiga*	‘the friend’ (fem.)
		[ˈkazɐ] + [ɐˈzu̯ɫ] → [ˈkazaˈzu̯ɫ]	*casa azul*	‘blue house’
		[ˈpagɐ] + [ɐˈkõ^{n}tɐ] → [ˈpagaˈkõ^{n}tɐ]	*paga a conta*	‘pay the bill!’
		[ˈɛrɐ] + [ɐˈli] → [ˈɛraˈli]	*era ali*	‘it was there’

Carvalho (2011) analyzes Portuguese vowel height using an element-based model, in which height is represented by combining different numbers of primitive “resonance elements” A. More elements A are used to represent lower vowels, and so a phonologically low vowel has two elements A. [ɐ] (equivalent in height by his analysis to [ʌ]) has a different representation from phonological /a/; [ɐ] has one A, while [a] has two As:

(29) Element-based representations of vowels (Carvalho 2011)

/i/ = {I}		/u/ = {U}
/e/ = {I,A}		/o/ = {U,A}
	/ʌ/ = {A}	
/ɛ/ = {I,A,A}		/ɔ/ = {U,A,A}
	/a/ = {A,A}	

Carvalho suggests that crasis can be explained by a sort of additive effect of A elements. When two instances of [ɐ] come together across a word boundary, each contains one element A. When these two combine, however, the vowel contains two A elements, and is thus equivalent to [a] = {A,A}. I have illustrated this in (30):

(30) Illustration of Carvalho’s (2011) analysis of crasis

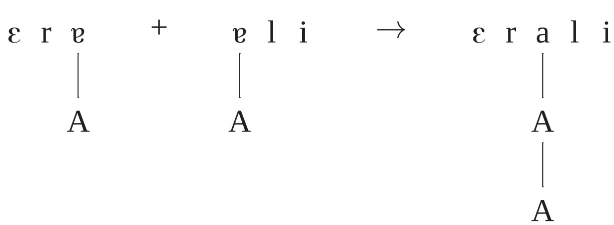

There are several problems with such an analysis. Because the privative elements of these kinds of representations are used to reflect phonological contrasts between categories, they are equivalent in some sense to contrastive features. If different allophones are specified with the same contrastive features as contrasting phonemes, then they cease to be merely allophones. This instead conflates the notions of phonological and phonetic vowel reduction. Because it is a principle of my model to not change the featural content of underlying phonemes positionally upon reduction if they do not undergo any neutralization, this analysis is incompatible as is. Even taking a weaker view of what it means to be a contrastive feature, this analysis requires privative height features or elements in some way to work, which, as mentioned in the preceding section, may not necessarily be compatible with a contrastive hierarchical approach.

Another problem with analyzing crasis as the combination of phonological elements is that, by operating on the representations in (29), which apparently refer to contrasts, it presupposes discrete categories. However, crasis does not actually seem to be a categorical process. According to Parkinson (1983:169), "contraction [crasis] is not an all-or-nothing process, which either happens or does not happen. Various degrees of contraction can be identified." His illustrations of these are given in (31):

(31) Variable coalescence in European Portuguese crasis (Parkinson 1983:169)

Maximum				Minimum
V	V	V‿V	VV	V V
[kazazul]	[kaza‿ᵄzul]	[kazɐ‿ɐzul]	[kazɐɐzul]	[kazɐˀɐzul]

While the way Carvalho uses phonological elements is incompatible with the notions of contrast and neutralization as I define them here, his analysis is insightful in the way that it uses an additive effect of privative elements. I propose a different analysis of crasis which uses such an additive effect. Recall first Barnes's (2006) explanation of the apparent distinction between moderate and radical reduction realizations of the neutralized /a/–/o/ pair in Russian unstressed syllables. This categorically defined pair (one non-terminal node "phoneme" in my model) has been described as having one of two realizations [ɐ] or [a] depending on prosodic position within a word.

In reality, however, the height of the unstressed /a/–/o/ pair depends not on position within a word, but on duration. Height is very predictably a function of duration, and is thus a phonetic property of the utterance, not a matter of categorical reductions in the phonology.

If the same non-contrastive duration-based variation in height realization holds for unstressed European Portuguese /a/ as for Russian /a/–/o/, then we would expect that crasis-affected /a/ should have longer duration. This is where the additive effect of privative elements becomes a useful mechanism: assuming that vowels in Portuguese underlyingly have one V timing slot, then the coalescence of two vowels across a word boundary can be analyzed as two identical adjacent melodies being interpreted as one single melody with with two timing slots, as illustrated in (32). Because timing slots are representations of duration, a vowel with two V slots would be phonetically spelled out as longer. If vowel height is a function of this duration, then a blocking of allophonic unstressed /a/ raising is exactly what we would expect in the coalescence of crasis.

(32)	V C V	+	V C V	→	V C V V C V
	/ɛ r a/		/a l i/		/ɛ r a l i/
	[ˈɛrɐ]		[ɐˈli]		[ˈɛɾaˈli]

The gradience of contraction in (31), then, can be said to follow from duration which varies due to whether or not the two V slots link to the same melody, and to the vowel's phonetic duration as a function of the degree to which the two adjacent words come together in a phonological phrase.

The timing-slot analysis of crasis I have proposed here makes several testable predictions. First, if the realized height of /a/ is a function of duration, then the [a] of forms which have undergone crasis should be measurably longer than a single unstressed [ɐ] of a form which has not. Second, at higher speech rates, even crasis vowels which have completely coalesced into a single segment articulatorily should be realized with raising, rather than as [a], provided that they have shorter duration.

Under at least a preliminary look, these predictions seem to hold. According to at least one corpus study of the relationship between vowel height and duration in European Portuguese (Gendrot and Adda-Decker 2007), shorter durations do indeed seem to be correlated with height of

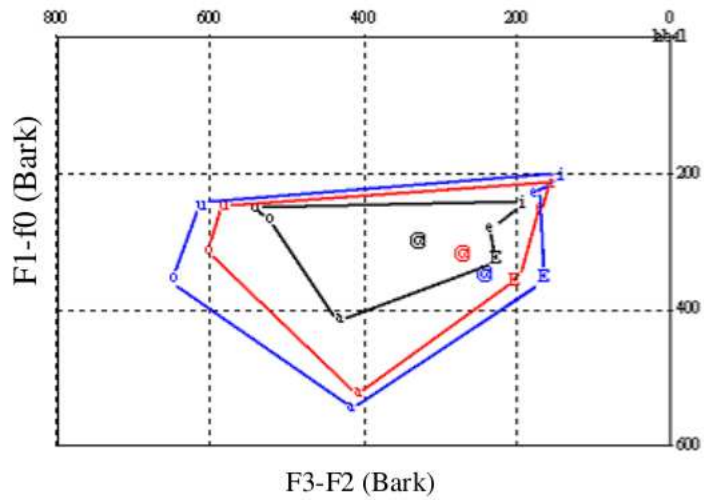

"From inside to outside (black [30–50ms], red [60–80], blue [90–110]). Male and female results are merged."

Figure 1: Portuguese vowel formants based on duration (Gendrot and Adda-Decker 2007)

/a/ as expected, as seen in Figure 1. This figure shows duration of vowels divided up into three different durational windows, and plotted on a single graph. The outermost polygon shows the average formant values of vowels with the longest durations, from 90 to 110 milliseconds, while the middle polygon shows the same for vowels from 60 to 80 milliseconds. The inner polygon reflects durations from 30 to 50 milliseconds. By looking at the three positions of the bottom point of these shapes, it can be clearly seen that /a/ is realized higher at shorter durations and lower at longer durations.

Unfortunately, there are some problems with the way the data is presented, for which Gendrot and Adda-Decker offer no explanation. Not all of the vowel phonemes of Portuguese seem to be represented in the figure. They do state in their text that Portuguese has a phoneme /ɔ/, but do not plot it. They do, however, plot the corresponding front vowel /ɛ/, which appears in the figure as "E". They also do not explain what is meant by the "@" that appears three times, colour-coded for duration along with the other symbols.

Another issue is that there is no specific mention made about the kinds of phonological environments in which the segments occur, such as whether or not they are stressed and whether they occur across word boundaries.[14] In order to make up for these shortcomings, future research will involve more focused instrumental investigation of duration and vowel height in European Portuguese crasis contexts.

4.4 Verbal morphology and the role of /a/

Most descriptions of Portuguese assume that verbs consist of a root, followed by a theme vowel, which is one of /a, e, i/. The root and theme vowel together constitute the verb stem, which is followed by a Tense–Aspect–Mood suffix, and then a Person–Number suffix. This is schematized in (33).

(33) Structure of the Portuguese Verb

	Verb Stem					
[Root	+	Theme V]	+	TAM Suffix	+	PN Suffix

The Person–Number suffixes can vary from tense to tense. The present indicative uses the Person–Number suffixes in (34). The second person plural forms are not provided because, as Mateus and d'Andrade (2000:73) note, this form is not common, being used only in a few dialects.

(34) Present indicative person and number suffixes (Mateus and d'Andrade 2000:74)

1st sing.	/o/	1st plur.	/mos/
2nd sing.	/s/	2nd plur.	(none)
3rd sing.	ø	3rd plur.	/N/ (nasal autosegment)

The present indicative is assumed not to have any Tense–Aspect–Mood morpheme, and so affixing the suffixes in (34) to verb stems gives forms as exemplified in (35).[15] Note that by this description, the theme vowel is deleted preceding a vowel-initial Person–Number suffix, as is the case with the first person singular forms.

[14] This is likely because Gendrot and Adda-Decker (2007) was a study of automatic formant measuring of annotated corpora from eight different languages, and so this level of detail was not readily attainable.

[15] In (35) and (37) below, the second "+" indicates the division between the theme vowel and the person suffix, which Mateus and d'Andrade (2000) notate as the stem edge "$]_{\mathrm{St}}$". I have changed this for readability.

(35) Present indicative forms (Mateus and d'Andrade 2000:74, adapted)

/a/ themes		/e/ themes		/i/ themes	
/fal+a+o/	[fálu]	/bat+e+o/	[bátu]	/part+i+o/	[pártu]
/fal+a+s/	[fálɐʃ]	/bat+e+s/	[bátɨʃ]	/part+i+s/	[pártɨʃ]
/fal+a+/	[fálɐ]	/bat+e+/	[bátɨ]	/part+i+/	[pártɨ]
/fal+a+mos/	[fɐlɐ́muʃ]	/bat+e+mos/	[bɐtémuʃ]	/part+i+mos/	[pɐrtímuʃ]
/fal+a+N/	[fálɐ̃w̃]	/bat+e+N/	[bátɐ̃j̃]	/part+i+N/	[pártɐ̃j̃]

There are a number of instances of surface allophony in (35) which should be pointed out. For example, underlying /s/ is regularly realized as [ʃ] in coda position. Vowels in unstressed syllables have, of course, different realizations than their underlying forms, following the patterns of phonological and phonetic vowel reduction described above. Finally, /a/ is realized as [ɐ] even in stressed position in the first person plural form 'we speak' [fɐlɐ́muʃ]. According to Mateus and d'Andrade (2000:19), it is a regular generalization of European Portuguese that stressed /a/ is allophonically raised to [ɐ] before nasals and palatals.

Stressed /a/ raising is not a problem for an analysis in which /a/ never undergoes any kind of neutralization; we simply need to add prenasal and prepalatal position to the list of of environments in which the raised allophone of /a/ is realized. However, the first person plural past perfect indicative forms of /a/ theme verbs show a systematic exception to prenasal /a/ raising, which gives rise to an apparent contrast between [a] and [ɐ]. The past perfect indicative, like the present indicative, is said not to have a Tense–Aspect–Mood suffix. Instead, the past perfect is encoded fusionally, using a separate set of Person–Number suffixes, given in (36). Note, however, that the underlying form of the first person plural suffix is identical to that of the present.

(36) Past perfect indicative person and number suffixes (Mateus and d'Andrade 2000:77)

1st sing.	/i/	1st plur.	/mos/
2nd sing.	/ste/	2nd plur.	(none)
3rd sing.	/u/	3rd plur.	/ɾaN/ (nasal autosegment)

When these suffixes are added to verb stems, the surface forms in (37) are found:

(37) Past perfect indicative forms (Mateus and d'Andrade 2000:77, adapted)

/a/ themes		/e/ themes		/i/ themes	
/fal+a+i/	[fɐlɐ́j]	/bat+e+i/	[bɐtí]	/part+i+i/	[pɐrtí]
/fal+a+ste/	[fɐláʃtɨ]	/bat+e+ste/	[bɐtɛ́ʃtɨ]	/part+i+ste/	[pɐrtíʃtɨ]
/fal+a+u/	[fɐló]	/bat+e+u/	[bɐtéw]	/part+i+u/	[pɐrtíw]
/fal+a+mos/	[fɐlámuʃ]	/bat+e+mos/	[bɐtémuʃ]	/part+i+mos/	[pɐrtímuʃ]
/fal+a+raN/	[fɐlárɐ̃w̃]	/bat+e+raN/	[bɐtérɐ̃w̃]	/part+i+raN/	[pɐrtírɐ̃w̃]

The forms in (35) and (37) therefore give a surface minimal pair, despite the fact that they have the same underlying forms as described by Mateus and d'Andrade (2000):

(38) a. /fal+a]$_{\mathrm{St}}$mos/ [fɐlɐ́muʃ] 'we speak'

b. /fal+a]$_{\mathrm{St}}$mos/ [fɐlámuʃ] 'we spoke'

This apparent contrast applies, as far as I can ascertain, to these forms in all /a/ theme verbs, yet Mateus and d'Andrade (2000) seem not to mention it at all. Redenbarger (1981) describes some authors as having posited two separate phonemes /a/ and /ɐ/ to account for the distinction in (38). Even Carvalho (2004), who offers a very insightful analysis of other verbal alternations, sets aside this behaviour of /a/:

> "I shall leave aside the /ɐ/ ∼ /a/ opposition that exists only before heterosyllabic nasals in central [European Portuguese] for a particular morphological purpose: *matamos* 'we kill' / *matámos* 'we killed'." (Carvalho 2004:14, note 2)

In the following section, I will extend Carvalho's (2004) analysis of Portuguese verbs in order to account for the [a]∼[ɐ] alternation using neither two separate underlying phonemes nor the kinds of morphophonological allophonic rules implied by the quote above.

4.5 Templates and autosegments in the Portuguese verb

Carvalho (2004) gives an analysis of certain root vowel alternations in Portuguese verb forms. Almost all Portuguese verbs with mid vowels in their roots have the lower mid vowels /ɛ/ or /ɔ/. In forms whose Person–Number suffixes start with a vowel, such as the first person singular ending in /-o/, the stem vowel surfaces with the height of the theme vowel. Thus /ɛ, ɔ/ stems with /e/ themes surface as /e, o/, and those with /i/ themes surface as /i, u/:

(39) Portuguese verb root "harmony" (Carvalho 2004:19, adapted)

a.	ɛ + a	→	[ɛ]	/lɛva/ + o	→	[lɛv]o
	ɔ + a	→	[ɔ]	/mɔra/ + o	→	[mɔr]o
b.	ɛ + e	→	[e]	/mɛte/ + o	→	[met]o
	ɔ + e	→	[o]	/kɔme/ + o	→	[kom]o
c.	ɛ + i	→	[i]	/fɛri/ + o	→	[fir]o
	ɔ + i	→	[u]	/tɔsi/ + o	→	[tus]o

Wetzels (1995, 2010) accounts for these alternations with a kind of vowel harmony that he calls "deletion-cum-spreading"; the height features of a theme vowel spread to the root vowel provided the theme vowel deletes. This deletion takes place when the suffix added to the verb stem is vowel-initial. However, given that the effect of spreading only occurs when, within Wetzels's proposal, deletion takes place, the data seem to call out for a kind of relinking analysis. Carvalho (2004) proposes instead that Portuguese verbs are based on autosegmental CV templates, where it is not necessary for each morpheme to occupy its own timing slot; morphemes can contain empty timing slots or floating melodies. Several morphemes then, including the theme vowels themselves and the first person singular suffix, consist solely of floating melodies:

(40) Morphemes lacking skeletal positions (Carvalho 2004:24)

a. The thematic vowels *a, e, i.*
b. The 1st person suffix *o,* and the subjunctive affixes *e, a.*

There is a restriction on the shapes of root templates, such that the last syllable contains only a single V timing slot. This will underlyingly have no vowel melody associated to it. By convention, melodies are associated to the CV template from right to left, and so in forms in which the only floating vowel is the theme vowel, its melody is associated to the final syllable, and the stem vowel is left unaffected. This is shown in (41) for the third person singular forms of the verbs *meter* 'to put’ and *comer* ‘to eat’, where there is a null Person–Number suffix.[16]

(41) Association of theme vowel melody (Carvalho 2004:24, adapted)

a. *mete* [ˈmɛt^{ə}] b. *come* [ˈkɔm^{ə}]

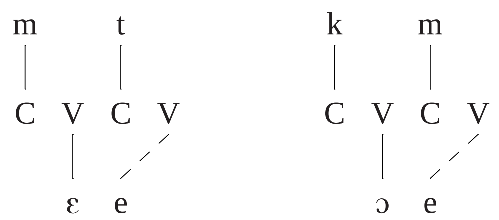

When a different morpheme, such as the first person singular suffix, leaves an extra vowel melody on the right edge of the underlying form, it first associates to the extra V timing slot on the verb stem. Because melodies associate from right to left, the floating theme vowel instead associates to the root vowel, spreading its height features and causing “harmony”:

(42) Stem vowel harmony (Carvalho 2004:24, adapted)

a. *meto* [ˈmet^{u}] b. *como* [ˈkom^{u}]

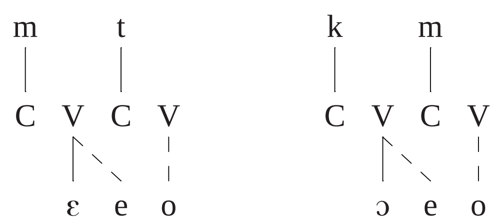

4.6 A unified analysis of the blocking of /a/ raising

We saw in §4.3 that crasis, the blocking of /a/ raising in unstressed syllables across word boundaries, seems to be triggered by the presence of two V timing slots associated with a single /a/ vowel melody. If being associated with two V slots is the phonological configuration which results in phonetic realization of /a/ as [a] in one environment in which its raised allophone [ɐ] is usually realized, then it is worth considering the effects that such representations may have in other raising environments.

[16]The superscript vowels in (41) and (42) are apparently meant to represent the very short duration in word-final position. Recall from the discussion of European Portuguese vowel reduction above that both of these vowels can optionally delete word-finally.

Recall that Mateus and d'Andrade (2000) posit a null morpheme for both the present and past perfect tenses, claiming that they are distinguished by different sets of Person–Number suffixes alone. This causes a problem because the first person plural suffixes in both tenses are assumed to be homophonously /-mos/, meaning identical underlying forms for the (in some dialects of European Portuguese, non-homophonous) forms of /a/ theme verbs fal[ɐ́]mos/fal[á]mos 'we speak/we spoke'. Given that there is clearly a distinction here, however, we should consider what may account for it.

Carvalho (2004) has given good reason to believe that the timing and melodic tiers can function and be represented autonomously of each other in Portuguese verbal morphology, and I will show that the behavior of /a/ theme verbs provides independent evidence for this mechanism. I do this by proposing that there is in fact a past perfect morpheme in European Portuguese, appearing in the Tense–Aspect–Mood position, which takes the phonological form of a floating V timing slot, with no associated melody:

(43) 'past perfect'

/ – V /

The underlying form of the root 'speak' follows exactly the CVCV verb template that Carvalho proposes, with no melody associated to the final V slot, but a floating melody for its theme vowel /a/, as in (44a). The first person plural suffix in (44b) (identical in the present and past perfect), as Carvalho (2004:27) describes it, does have a specified CV tier.

(44) a. 'speak'

f l

/ C V C V /

a a

b. 'first person plural'

m s

/ – C V C /

o

When the present tense is formed, the Person–Number suffix from (44b) is attached to the right edge of the stem. Floating melodies are then associated from right to left, and so the /a/ theme vowel associates to the empty V slot, as in (45). Because the /a/ only occupies a single timing slot, it systematically undergoes allophonic raising to [ɐ] before the nasal /m/ which it precedes.

(45) [fɐlɐ́muʃ] ‘we speak’

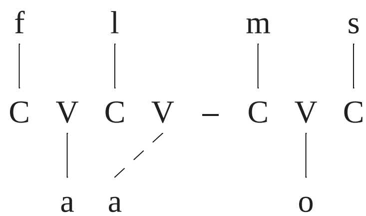

When the past perfect is formed, both the Tense–Aspect–Mood suffix in (43) and the first person plural suffix are added to the stem. Floating melodies are then associated from right to left. However, as seen in (46), the floating theme vowel melody /a/ has two empty V slots to which it can associate, and so it associates to both of them. Although this /a/ precedes a nasal, and so would be expected to undergo allophonic raising to [ɐ], the fact that it is dominated by two V slots prevents this raising from taking place, as it does with crasis.

(46) [fɐlámuʃ] ‘we spoke’

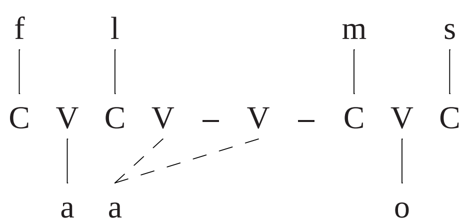

As with the analysis of crasis discussed in §4.3, positing two V slots on the non-raised prenasal /a/ in the past tense forms makes several testable predictions, which can be more closely investigated in future research: if duration is correlated with height, then we should expect the [a] in the past tense form in (46) to be measurably longer than the stressed [ɐ] in (45), and furthermore that the differences will become less distinct or possibly eliminated eliminated at faster speech rates.

It should be noted as well that moving towards an analysis where the past perfect morpheme is only a V slot to account for the first person plural forms has minimal consequences for the other forms. According to Carvalho (2004:27), the Person–Number morphemes for the past perfect are among the suffixes in the language that have their own CV structure, rather than consisting of only floating melodies. They are also fusional, containing both Person–Number and Tense–Aspect–Mood information:

(47) Past perfect Person–Number as CV suffixes (Carvalho 2004:27, adapted)

i	‘1st person singular’ + ‘perfect’
s]te	‘2nd person singular’ + ‘perfect’
u	‘3rd person singular’ + ‘perfect’
ram	‘3rd person plural’ + ‘perfect’

Instead, the first person singular and third person singular past perfect suffixes would also consist only of floating melodies, as their V slot is now provided by the past perfect morpheme. For example, with the past perfect first person singular of ‘speak’, the floating /i/ first associates with the V slot of the past perfect, followed by the floating /a/ theme associating with the extra timing slot on the root. Together, these form a diphthong [ɐj], where the /i/ glides to [j], and the /a/, despite being stressed, allophonically raises to [ɐ] because it precedes a palatal:

(48) [falɐ́j] ‘I spoke’

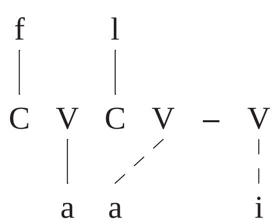

The second person singular and third person plural forms then work similarly to the first person plural; they have their own CV and melodic representations. In these forms, the theme vowel melody associates to both the extra V slot on the root and the V slot of the past perfect morpheme, as shown in (49). However, since the environments in which the stressed /a/ occurs (before [ʃ] and [r] respectively) do not condition allophonic raising, there is no place in which an apparent surface contrast arises due to the blocking of raising.

(49) [fɐláʃtɨ] ‘you spoke’

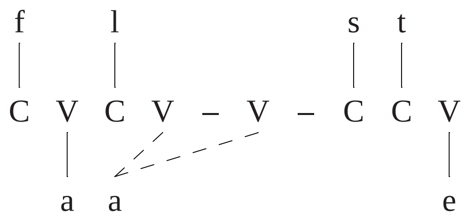

We may ask at this point how the past perfect morpheme came to be represented as merely a V timing slot. The diachronic facts offer some additional corroboration for such an analysis. Diachronically we would expect additional length to have originated through compensatory lengthening, i.e. through the loss of a syllable. This is precisely what we find: as seen in (50), the first and third person singular, and (crucially) the first person plural forms did indeed lose a syllable in the transition from Latin to Modern Portuguese.[17]

(50) Diachrony of Portuguese past perfect suffixes (Williams 1962:193)

Classical Latin	Vulgar Latin	Portuguese
-āvī	*-ái*	*-ei*
-āstī	*-ásti*	*-aste*
-āvit	*-áut*	*-ou*
-āvĭmŭs	*-ámus*	*-amos* [amuʃ]
-āstĭs	*-ástes*	*-astes*
-ārunt	*-árunt*	*-arom* > *-aram*

Once the extra V slots entered into speakers’ representations for some of the past perfect forms, they could have been extended to the rest of the past tense forms, where the V was reanalyzed as the past morpheme itself. Williams (1962) actually offers a somewhat similar but incompatible explanation for the [a] appearing in the past tense form in Modern Portuguese. He says that the non-raised [a] of the past tense form (in dialects which have this distinction) was due to analogy with the second person singular and the second and third person plural forms, which all have a stressed [a].

[17] The present tense third person plural ending in Classical Latin, like the past tense form, contained a long vowel: *-āmŭs*, which became Modern Portuguese *-amos* [ɐmuʃ] (Williams 1962:188). While adding length (which was at some point lost as a surface distinction) to the picture may complicate an explanation based on compensatory lengthening, it is clear that the present tense form lost merely the length distinction, while the past perfect form lost both length and a full syllable, so we should expect that at some point an additional timing slot was present in the representations of the past perfect form that was not present in the present form.

This is similar in that it involves some paradigm leveling, but incompatible with contrast as I define it in this paper. Saying that one vowel can be used analogically in place of another as a phonological distinction implies that there are, or at least were at some point, two separate phonemes /a/ and /ɐ/. However, there seems to be no evidence for this, and the two sounds appear to be in contrast with each other only in a limited set of forms, namely the first person plural forms of the /a/ theme verbs.

I have shown in this case study of European Portuguese /a/ that, despite this apparent contrast, it is not necessary to have two different phonemes for the different realizations of /a/. In contexts in which /a/ is expected to be realized as [ɐ] but surfaces instead as [a], there is a common factor: the /a/ melody is associated with two V timing slots, and there is independent evidence for the existence of these timing slots. In the case of crasis, the two timing slots can be seen to come together even in a gradient fashion, and are necessarily underlying in the forms as they would be pronounced in isolation. With regards to the role of an extra V slot in the past tense verb forms, there are independent reasons to believe that timing slots and melodies interact independently of each other in Portuguese verbal morphophonology. Finally, the multiple V slot analysis of /a/ non-raising makes testable predictions about the duration of /a/ when realized as [a] as opposed to [ɐ] in the same context, which can be investigated more closely in future research.[18]

5 Theoretical issues in consonant neutralization

In the preceding sections I have shown how positional neutralization in phonological vowel reduction can be represented in a contrastive hierarchical system through the use of non-terminal nodes, and I have addressed ways of analyzing apparent surface contrasts which do not involve neutralization. The use of non-terminal nodes of the contrastive hierarchy to represent true neutralization

[18] This analysis also seems to make a prediction about the duration of other vowels that appear with two V slots in past perfect forms, such as in non-/a/ theme verbs. Whether the duration actually varies along the same lines here should be apparent under phonetic investigation. It does seem to be the case, however, that the vowel /a/'s height is affected more greatly by durational differences than other vowels, as seen in Table 1, and so this may be some special property of the lower vowels within a given system. I leave investigation of this to future research.

is useful for a number of reasons, and makes a number of predictions about the ways that different neutralizations can relate to each other, both phonologically and phonetically. Because the non-terminal node model of neutralization is intended to deal with all phonological neutralizations, not only vowel reduction, this section will address some theoretical advantages and concerns arising from the application of the model to consonantal contrasts.

5.1 Neutralization of multiple features

In some cases, a set of more than two terminal segments may neutralize by more than one feature depending on context. The non-terminal node model of neutralization which I have laid out in the preceding sections makes some predictions about the ways that multiple neutralizations can coexist within the same system. Specifically, all segments which neutralize for a particular feature must also neutralize for all features ranked below it on the hierarchy. I will illustrate this with data from Bulgarian.

The vast majority of consonants in Bulgarian occur in palatalized/non-palatalized pairs (see Table 1, where palatalization is indicated with a superscript ^{j} on the consonant). Furthermore, almost all of the obstruents contrast in voiced/voiceless pairs. However, there are a number of contexts in which voicing and palatalization are neutralized.

According to Scatton (1975:viii), palatalization is only contrastive before back vowels. It is not contrastive word-finally, in clusters before obstruents, liquids, or nasals, or before front vowels. He claims that the “hard” (non-palatalized) variants occur in codas (before consonants and word-finally), and that velars appear as hard before front vowels. The “soft” (palatalized) variants of liquids, nasals, and non-velar obstruents appear before front vowels. In other words, consonants are only contrastive for palatalization before non-front vowels.

Voicing is not neutralized in quite as many positions; Scatton (1975:viii–x) states that voicing is contrastive before all sonorants. Thus, unlike palatalization, voicing *is* contrastive in coda position before a sonorant consonant, and it is contrastive before all vowels, not just before non-front vowels (like palatalization). At the end of a word, he claims that obstruents occur in their voiceless

			labial	alveodental	alveopalatal	palatal	velar
obstruents	stops	voiceless	p p^j	t t^j		ḱ	k
		voiced	b b^j	d d^j		ǵ	g
	fricatives	voiceless	f f^j	s s^j	š		x
		voiced	v v^j	z z^j	ž		
	affricates	voiceless		c c^j	č		
		voiced			ǯ		
sonorants	nasals		m m^j	n n^j			
	liquids	lateral		l l^j			
		trill		r r^j			
	glide			j			

Table 1: Bulgarian consonant inventory (Scatton 1984:59, adapted).

variants. In clusters, voiced variants occur before voiced obstruents, and voiceless variants occur before voiceless consonants.

Much as for vowel reduction, we cannot know the correspondences between stressed (non-reduced) and unstressed (reduced) vowels without examining forms of the same morpheme with stress on different syllables. Likewise, finding the correspondences of consonantal neutralization requires, for example, looking at different forms of words in which consonants, often at morpheme edges, occur in different contexts. For example, there are some consonant-final noun stems ending in palatalized consonants, such as /kon^j/ 'horse'. In the singular, without a suffix, the underlying palatalization is neutralized, and the stem appears to be like /ston/ 'moan'. However, when suffixed, the contrast becomes apparent:

(51) Word-final neutralization of palatalization (Zec 2002:240–241, adapted)

a. /kon^j/ $\rightarrow$ [kon] 'horse' /kon^j-ât/ $\rightarrow$ [kon^{j}ət] 'the horse'

b. /ston/ $\rightarrow$ [ston] 'moan' /ston-ât/ $\rightarrow$ [stonət] 'the moan'

Similarly, stems showing word-final neutralization of voicing can be seen, as in (52a), where the underlyingly voiced stem-final consonant surfaces as voiceless word-finally but voiced intervocalically, as opposed to (52b), where the underlyingly voiceless stem-final consonant always surfaces as voiceless.

(52) Word-final neutralization of voicing (Aronson 1968:35, adapted)

a. /nož/ → [noš] ‘knife’ /nož-ove/ → [nož-ove] ‘knives’
b. /svat/ → [svat] ‘matchmaker’ /svatove/ → [svatove] ‘matchmakers’

When comparing the contrasts of voicing and palatalization, we can see in (53) that the contexts in which voicing is neutralized are a subset of the contexts in which palatalization is neutralized. In other words, it is possible to neutralize palatalization without neutralizing voicing, but it is not possible to neutralize voicing without neutralizing palatalization.

(53) Positional neutralization of palatalization and voicing in Bulgarian

Palatalization	Voicing
Word-finally	Word-finally
In all clusters	In clusters before obstruents
Before front vowels	

We can therefore conclude that the terminal contrast is for palatalization, represented here with the feature [±front], and that voicing is ranked over that, as in the tree in (54). In contexts in which palatalization is neutralized, rules change nodes 4 and 5 to node 2, and change nodes 6 and 7 to node 3. This works in exactly the same way as it does for vowel reduction in contexts which neutralize for height features. In contexts in which voicing is neutralized, and thus palatalization is also obligatorily neutralized, nodes 4, 5, 6, and 7 are all rewritten to node 1. Thus, in contexts which neutralize voicing, obstruents are only contrastive for place and manner.

(54) Contrastive hierarchy for Bulgarian alveolar stops

- (alveolar stop)$_1$
  - [+voice]$_2$
    - [+front]$_4$ /d^j/
    - [–front]$_5$ /d/
  - [–voice]$_3$
    - [+front]$_6$ /t^j/
    - [–front]$_7$ /t/

We saw in the preceding sections that non-terminal nodes (such as the Russian /a/–/o/ node in §3.2) can have their own allophony, because they have status as members of the inventory, just like terminal nodes. Since phonetic realization is not a diagnostic for contrastive phonemic status, when Scatton (1975) says that all consonants are realized as non-palatalized before front vowels, this does not mean that they are contrastively [–front] in these positions, but merely that it is a non-palatalized allophone of the non-terminal node above [±front] that is realized. Indeed, allophony for palatalization seems to be more fine grained than this. Scatton (1984:64–65) notes that "in these environments they moderately assimilate to the tonal qualities of the vowel: before /i/ they show weak i-tonality, before /e/ weak e-tonality." Furthermore, in the Eastern dialects of Bulgarian, it is common for (palatalization contrast neutralized) consonants to actually be pronounced as palatalized before /i/ and /e/, rather than non-palatalized. Using only terminal nodes, it would be necessary to say that some dialects use different terminal nodes in positions of neutralization, but within my model, it is instead the case that non-terminal nodes (which are not contrastive for the relevant feature) simply have different degrees of palatalized allophony, caused by varying degrees to which the frontness of following vowels is allowed to show coarticulation on the consonant.

The same can be said for voicing. Voicing (along with palatalization) is neutralized in clusters before obstruents, but is phonetically realized with the voicing of the following consonant. Rather than changing its voicing feature phonologically to that of the following segment, obstruents in clusters are represented as the non-terminal node above the voicing contrast, which merely has allophonically voiced and voiceless variants. In the case of clusters which never undergo alternations because they are not across morpheme boundaries, the segments which do not show a contrast may be represented underlyingly with the non-terminal node rather than having a specification for voicing. This yields greater representational economy by requiring the specification of fewer features.

By using a non-terminal node approach to the neutralization of features, we are afforded a principled and restrictive way of representing the neutralization of multiple features. The system here predicts that in cases of true phonological neutralization of two different features for some set of terminal contrasts, the contexts in which the lower feature in the hierarchy neutralizes will be a superset of the contexts in which the higher feature neutralizes. To take the tree in (54) for example, the model presented here would not predict that a position could neutralize palatalization without neutralizing voicing.

5.2 Underlying non-terminal nodes: Three-way contrasts?

In the model that I have proposed in this paper, it is possible to use non-terminal nodes of the contrastive hierarchy underlyingly. So far, this possibility has only been exploited to make underlying representations more economical and less arbitrary; in cases where a language learner encounters a segment in neutralized position, and there is no evidence, either from the same morpheme appearing in a context in which the segment is not neutralized, or from some sort of phonological process, a non-terminal node may be posited underlyingly. In this way, one of the possible terminal nodes need not be chosen arbitrarily, only to be changed by rule every single time that the form is used. I would now like to consider briefly whether it is possible to use underlying non-terminal nodes in order to reflect different kinds of phonological activity. In other words, can a non-terminal node and the two terminal nodes it dominates together provide a three-way phonological contrast?

Let us first consider what such a three-way situation might look like. In a typical voicing system, such as we saw above for Bulgarian, obstruents may be contrastive for a feature [±voice], each having its own phonetic realization. In certain positions, such as word-finally or in clusters before other obstruents, this contrast may be neutralized, meaning that its phonetic voicing is predictable from context, and there can be no opposition between voiced and voiceless. If neutralization is analyzed as a change from a phoneme which is contrastive for voicing to an archiphoneme (a non-terminal node) which is not contrastive for voicing, then phonemes specified for both [+voice] and [–voice] lose their contrastive specifications for voicing, causing neutralization. Evidence for neutralization, and thus the use of a non-terminal node, comes from the inability of a position to show one or the other of the lower nodes. But let us consider the following data from Turkish, given by Inkelas (1996:3):[19]

(55) Types of voicing contrasts in Turkish (Inkelas 1996:3, adapted)

a. Alternating root-final plosive:

kanat	'wing'	kanad-ɯ	'wing-acc'
kanat-lar	'wing-pl'	kanad-ɯm	'wing-1sg.poss'

b. Non-alternating voiceless plosive:

sanat	'art'	sanat-ɯ	'art-acc'
sanat-lar	'art-pl'	sanat-ɯm	'art-1sg.poss'

c. Non-alternating voiced plosive:

etyd	'study'	etyd-y	'study-acc'
etyd-ler	'study-pl'	etyd-ym	'study-1sg.poss'

Turkish shows a surface two-way distinction in voicing on obstruent stops, but the distinction is neutralized in coda position, as seen in (55a) and (55b). Thus the root 'wing' appears with a final [d] when it serves as an onset, but as [t] when it is a coda. The final consonant in the root 'art', however, is always voiceless, whether it is in onset or coda position. This sounds very much like word-final voicing neutralization in Bulgarian, until examples like that in (55c) are considered. Due to a number of recent loanwords entering the language, there are lexical items which have non-alternating *voiced* stops, which always appear as voiced, whether in coda or onset position. This means that only a subset of stops in Turkish have positionally predictable voicing; the rest are either always voiceless or always voiced. The question is: how should these three classes of stops be represented?

If we were concerned only with those in (55b) and (55c), we would think that /t/ and /d/ were

[19] I have changed the transcription of the vowels somewhat from Inkelas's (1996) original, but this has no bearing on the data discussed herein.

fully specified for [–voice] and [+voice] respectively, and that there was no neutralization at all. It is only the alternating stop in (55a), which I will refer to as /D/, that appears to undergo neutralization of its surface voicing. If the property of having predictable voicing is represented by having no contrastively specified voiced feature, then such three-way specifications are already predicted by a model allowing the use of non-terminal nodes underlyingly. They can be derived from the tree in (56):

(56) Voicing hierarchy for Turkish

(alveolar stop)$_1$

[+voice]$_2$ [–voice]$_3$
/d/ /t/

Here the terminal nodes 2 and 3 correspond to non-alternating /d/ and /t/, while the non-terminal node 1, not specified at all for [$\pm$voice], is free to vary positionally without violating its contrastive specifications or changing its phonemic identity, corresponding to /D/. This captures what Inkelas calls “archiphonemic underspecification”.

Such use of non-terminal nodes can also provide a reinterpretation of mixed voicing systems as analyzed by Avery (1996). Avery’s claim is that cross-linguistically, there are three different kinds of voicing systems, cast in privative feature geometric terms in (57).

(57) Representations of the voiceless-voiced opposition Avery (1996:126)

a. LV languages		b. SV languages		c. CV languages	
p	b	p	b	p	b
R	R	R	R	R	R
Lar	Lar		SV	Lar	
	[voice]				

(57a) shows the representations used in laryngeal voicing (LV) languages. Here both voiced and voiceless obstruents have a laryngeal node, and the voiced member is marked with a privative [voice] feature. The representations in (57b) are used in sonorant voice (SV) languages, where the voiceless member has only a bare root node, and the voiced member is marked with an empty SV node. Finally, in contextual voice (CV) languages, shown in (57c), the voiceless member has a bare laryngeal node, and it is the *voiced* member of the opposition which is unmarked, having only a bare root node.

The main place that Avery looks for evidence for the different kinds of systems in (57) is in the behaviour of voicing in codas. He assumes that various processes are in play to bring representations into conformity with the Laryngeal Condition (LC) in (58), a well-formedness constraint which says that all codas must consist of only a bare laryngeal node.

(58) Laryngeal Condition (LC) (Avery 1996:127)

$$\begin{array}{c} C]_{\sigma} \\ | \\ \text{Lar} \end{array}$$

The LC compels the application of processes delinking the [voice] feature on the voiced members in LV languages (57a), causing coda devoicing, and of adding a laryngeal node (but not [voice] feature) to the bare root node of the voiced member of CV languages (57c). However, segments with sonorant voicing are not affected by these processes, due to another constraint which prohibits an SV node and a laryngeal node to be present on the same segment:

(59) Lar-SV Constraint (Avery 1996:128)

$$\begin{array}{ccc} & {}^{*}\text{R} & \\ / & & \backslash \\ \text{Lar} & & \text{SV} \end{array}$$

Avery analyzes the Turkish voiceless stops and alternating voiced stops, seen in (55a, b), as fitting into a typical CV system as in (57c). Fricatives, however, are assumed to fit into an SV system, as in (57b), because voiced fricatives in Turkish do not devoice in coda position. The non-alternating voiced stops, which do not devoice in coda position, and therefore pattern with the voiced fricatives, are said to have an SV node, rather than a laryngeal node. In this way, Turkish shows a mixed voicing system. Avery's complete voicing specifications for obstruents, then as are in (60):

(60) Voicing specifications of Turkish obstruents (Avery 1996:150)

a. Alternating consonant (stops only)	b. Non-alternating voiced consonant (stops and fricatives)	c. Non-alternating voiceless consonant (stops and fricatives)
R	R	R
	|	|
	SV	Lar

If we attempt to convert such a three-way distinction into a contrastive hierarchy, where each of the three kinds of representations has its own terminal node, we can see that there are two possible trees:[20]

(61) Possible hierarchies for the /t, d, D/ opposition with Avery's (1996) specifications

a. (coronal stop)
- [+Lar] /t/
- [–Lar]
  - [+SV] /d/
  - [–SV] /D/

b. (coronal stop)
- [+SV] /d/
- [–SV]
  - [+Lar] /t/
  - [–Lar] /D/

The problem with the representations in (61) is that the relationship between /t/ and /d/ cannot be a minimal terminal opposition. Thus it cannot be analogous to the relationship between /s/ and /z/ in (62), with which it is supposed to pattern:

(62) Contrastive hierarchy for Turkish fricatives with Avery's (1996) specifications

(coronal fricative)
- [+SV] /z/
- [–SV] /s/

[20] I have converted the privative presence/absence of Lar and SV nodes into binary features here.

If we allow non-terminal nodes in underlying representations to be used contrastively, however, then the relationships between /s/ and /z/ is exactly analogous to that between /t/ and /d/. /D/ is not an additional terminal node disrupting these relationships, but rather a non-terminal node predicted to have exactly the kinds of specifications that Avery would require it to have, as seen in (63). The difference between the /s, z/ and /t, d/ opposition in Turkish, then, is not whether or not one allows a mixed voicing system, but rather whether or not one allows the contrastive use of non-terminal nodes.

(63) Hierarchies for analogous stop and fricative voicing systems

a. /D/

[+voice] [–voice]
/d/ /t/

b. (coronal fricative)

[+voice] [–voice]
/z/ /s/

Using the representations in (63a) underlyingly to reflect the voicing behaviour found in Turkish, no neutralization process is actually necessary at all, but merely an allophonic rule which predictably changes the surface voicing of the non-terminal node /D/. Before non-alternating /d/ was introduced into the system, however, there would only have been a two-way underlying contrast, between /t/ and /d/. In contexts of neutralization, both of the terminal nodes /t/ and /d/ would be changed by a phonological process to /D/. The addition of non-alternating /d/ to the system changed not the kinds of representations that could be interpreted, but rather the way in which underlying representations are handled. Instead of neutralizing phonological /t/–/d/ positionally, the positions in which there are surface alternations were reanalyzed to contain the underlying allophonically voiced segment, rather than acquiring a new terminal node. It was not the contrastive structure and available feature combinations of the inventory that changed, but rather the rules that can act upon it to cause neutralization of underlying contrasts.

Cases such as this, involving apparent use of non-terminal nodes to represent three-way distinctions in phonological activity with underlying representations, need to be investigated further in future research. Non-terminal nodes may prove useful in a more general theory of underspecification, because they provide a more restricted and non-arbitrary set of possible underspecified segments; with such a system, it is not possible to underspecify for any feature which is active in a given language, but rather only for features directly dominating the nodes which are contrastive in a given position, as determined by the contrastive hierarchy for the language. This restriction would not be afforded if new underspecified terminal nodes could be arbitrarily incorporated into the hierarchy.

6 Conclusion

In this paper, I have argued for a contrastive hierarchical model of positional neutralization in which non-terminal nodes of contrastive trees are available to the phonology as interpretable phonemes for the representation of segments in neutralized positions. This approach is held to be superior both conceptually and empirically to a literal subset approach, in which neutralized segments are represented by a subset of the possible terminal contrasts of the language, and identified based on their phonetic realization. The conceptual advantage is that correspondences between the full and reduced set of phonemes follows from a restriction of contrasts by disallowing any contrastive values of a given feature, compelling the use of a non-terminal node, rather than requiring some specification of a particular contrastive feature in a position in which that feature is not in opposition with anything. The empirical advantage follows directly from the representations used in the non-terminal node approach, in that the phonetic realizations of neutralized segments are predicted by implementing the feature specifications of the relevant non-terminal nodes rather than terminal ones.

This model allows a terminal and non-terminal node to have similar or identical phonetic realizations, because the diagnostic for using a non-terminal node is a positional inability to realize a contrast, not phonetic realization in that position. A corollary of this is that a single node can have positional allophones that sound very different, without altering the number of contrasts in any position. The realizations of these allophones can be determined by independently motivated phonological conditions, which can give rise to apparent surface contrasts, as I illustrated in §4 with a case study of European Portuguese. The analysis I developed there makes several testable predictions, which will be the focus of future research. Further work on positional neutralization will investigate several issues. The analysis I gave of Bulgarian vowel reduction in §3.1 captures the implicational hierarchy of neutralizing vowel pairs based on apparent markedness of feature values. By applying this model to more such implicational hierarchies with markedness in mind, it will be possible to determine if this result holds or was obtained merely by coincidence. A more intriguing and potentially drastic matter to consider is the possibility of using non-terminal nodes to represent underlying three-way contrasts, as discussed in §5.2.

References

Aronson, Howard I. 1968. *Bulgarian inflectional morphophonology*. The Hague and Paris: Mouton.

Avery, Peter. 1996. The representation of voicing contrasts. Doctoral Dissertation, University of Toronto.

Barbosa, Plínio A, and Eleonora Cavalcante Albano. 2004. Brazilian Portuguese. *Journal of the International Phonetics Association* 34:227–232.

Barnes, Johnathan. 2006. *Strength and weakness at the interface: Positional neutralization in phonetics and phonology*. Berlin and New York: Mouton de Gruyter.

Blevins, Juliette. 2004. *Evolutionary Phonology*. Cambridge: Cambridge University Press.

Carvalho, Joaquim Brandão de. 2004. Templatic morphology in the Portuguese verb. In *Nouveaux départs en phonologie*, ed. Trudel Meisenburg and Maria Selig, 13–32. Tübigen: Gunter Narr Verlag.

Carvalho, Joaquim Brandão de. 2011. Contrastive hierarchies, privative features, and Portuguese vowels. *Revista de estudos linguisticos da Universidade do Porto* 6:51–66.

Chomsky, Noam, and Morris Halle. 1968. *The sound pattern of English*. New York: Harper & Row.

Clements, Nick. 1991. Vowel height assimilation in Bantu languages. *Working Papers of the Cornell Phonetics Laboratory* 5:37–76.

Crosswhite, Katherine. 1999. Vowel reduction in Optimality Theory. Doctoral Dissertation, UCLA.

Crosswhite, Katherine. 2001. *Vowel reduction in Optimality Theory*. New York: Routledge.

Davidsen-Nielsen, Niels. 1978. *Neutralization and the archiphoneme: Two phonological concepts and their history*. Copenhagen: Akademisk Forlag.

Dresher, B. Elan. 2009. *The contrastive hierarchy in phonology*. Cambridge: Cambridge University Press.

Dyck, Carrie Joan. 1995. Constraining the phonology–phonetics interface: With exemplification from Spanish and Italian dialects. Doctoral Dissertation, University of Toronto.

Gendrot, Cédric, and Martine Adda-Decker. 2007. Impact and duration of vowel inventory size on formant values of oral vowels: An automated formant analysis from eight languages. *International Conference of Phonetic Sciences* 16:1417–1420.

Hall, Daniel Currie. 2011. Phonological contrast and its phonetic enhancement: Dispersedness without dispersion. *Phonology* 28:1–54.

Harvey, Christopher. 2012. Contrastive shift in Ob-Ugric vowel systems. Ms., University of Toronto.

Inkelas, Sharon. 1996. Archiphonemic underspecification: An optimization approach to the phonological description of morphemes. Ms., UC Berkley.

Iosad, Pavel. 2012. Vowel reduction in Russian: No phonetics in phonology. *Journal of Linguistics*. 48:521–571.

Mateus, Maria Helena, and Ernesto d'Andrade. 2000. *The phonology of Portuguese*. Oxford: Oxford University Press.

Oxford, Will. 2012. Patterns of contrast in phonological change: Evidence from Algonquian vowel systems. Ms., University of Toronto.

Parkinson, Stephen. 1983. Portuguese nasal vowels as phonological diphthongs. *Lingua* 61:157–177.

Prince, Alan, and Paul Smolensky. 2004. *Optimality Theory: Constraint interaction in generative grammar*. Oxford: Blackwell Publishing.

Redenbarger, Wayne J. 1981. *Articulator features and Portuguese vowel height*. Cambridge, Massachusetts: Harvard Studies in Romance Languages.

Scatton, Ernest A. 1975. *Bulgarian phonology*. Cambridge, MA: Slavica Publishers, Inc.

Scatton, Ernest A. 1984. *A reference grammar of modern Bulgarian*. Columbus, OH: Slavica Publishers, Inc.

Tilkov, Dimitar. 2009. *Le vocalisme bulgare*. Paris: Société de Linguistique de Paris.

Timberlake, Alan. 2004. *A reference grammar of Russian*. Cambridge: Cambridge University Press.

Trubetzkoy, N. S. 1969. *Principles of phonology*. Berkeley and Los Angeles: University of California Press.

Wetzels, Leo. 1995. Mid-vowel alternations in the Brazilian Portuguese verb. *Phonology* 12:281–304.

Wetzels, Leo. 2010. The representation of vowel height and vowel height neutralization in Brazilian Portuguese (Southern Dialects). In *Tones and features: Phonetic and phonological perspectives*, ed. John A. Goldsmith, Elizabeth Hume, and Leo Wetzels, 331–359. Berlin and Boston: De Gruyter Mouton.

Williams, Edwin B. 1962. *From Latin to Portuguese*. Philadelphia: University of Pennsylvania Press.

Zec, Draga. 2002. Constraints on multiple feature occurrence. In *Resolving conflicts in grammars: Optimality Theory in syntax, morphology, and phonology*, ed. Gisbert Fanselow and Caroline Féry, 237–264. Hamburg: Helmut Buske Verlag.