January 13, 2008. Slightly revised version to appear in *Natural Language and Linguistic Theory*.

WEIGHTED CONSTRAINTS AND GRADIENT RESTRICTIONS ON PLACE CO-OCCURRENCE IN MUNA AND ARABIC*

Andries W. Coetzee, University of Michigan

Joe Pater, University of Massachusetts, Amherst

Abstract. This paper documents a restriction against the co-occurrence of homorganic consonants in the root morphemes of Muna, a western Austronesian language, and compares the Muna pattern with the much-studied similar pattern in Arabic. As in Arabic, the restriction applies gradiently: its force depends on the place of articulation of the consonants involved, and on whether the homorganic consonants are similar in terms of other features. Muna differs from Arabic in the relative strengths of these other features in affecting co-occurrence rates of homorganic consonants. Along with the descriptions of these patterns, this paper presents phonological analyses of the Muna and Arabic patterns in terms of weighted constraints, as in Harmonic Grammar. This account uses a gradual learning algorithm that acquires weights that reflect the relative frequency of different sequence types in the two languages. The resulting grammars assign the sequences acceptability scores that correlate with a measure of their attestedness in the lexicon. This application of Harmonic Grammar illustrates its ability to capture both gradient and categorical patterns.

1. Introduction

The best-known restriction against homorganic consonants is the one that obtains between consonants in the Arabic verbal root (Greenberg 1950; McCarthy 1988, 1994; Pierrehumbert 1993; Padgett 1995; Frisch *et al.* 2004). The main interest of the Arabic pattern is that the strength of the restriction depends on several factors. One factor is the place of articulation of the consonants involved. Table 1 presents the lexical statistics for pairs of consonants at each of the four places of articulation: Labial, Coronal, Dorsal, and Pharyngeal. Only non-identical consonants are included, as we take the restriction against identical consonants, which is nearly absolute, to be a separate constraint (as in McCarthy 1994; cf. Frisch *et al.* 2004; see section 2 on Muna).

* Acknowledgements: We would like to especially thank René van den Berg, Stefan Frisch, and John McCarthy for sharing their data with us, and for providing helpful answers to our questions. In its long period of gestation, this paper also benefited tremendously from the comments of many others, whose names cannot all be listed due to space and memory limitations. We received useful written comments on earlier versions of the paper from Shigeto Kawahara and James Kirby, and especially from the Associate Editor and the anonymous reviewers. We gained much from discussions with Adam Albright, Jean-Pierre Angoujard, Arto Anttila, Pam Beddor, Diamandis Gafos, Gillian Gallagher, Bruce Hayes, René Kager, Michael Marlo, John McCarthy, Hugo Quené, Paul Smolensky and Colin Wilson. We received further useful feedback from participants at conferences and workshops where portions of this paper were presented, including the Manchester Phonology Meeting and the Seoul International Conference on Phonology, and at Indiana University, New York University, Radboud Univeristy, the University of Chicago, the University of Edinburgh, the University of Massachusetts, Amherst, the University of Michigan, the University of Nantes, and Utrecht University. This research was funded in part by the Nederlandse Organisatsie voor Wetenschappelijk Onderzoek under NWO grant #B 30-657.

Pharyngeal			Dorsal			Coronal			Labial		
ʕ ʔ ħ χ h ʁ			k q g			l r n Θ ð s z sˤ zˤ ʃ t d tˤ dˤ			b f m		
O	E	O/E	O	E	O/E	O	E	O/E	O	E	O/E
10	129.73	**0.08**	1	37.24	**0.03**	750	959.98	**0.78**	0	89.53	**0.00**

Table 1: Sequences of non-identical homorganic consonants in Arabic verbal roots

‘Observed’ (O) is the total number of sequences of two consonants with that place of articulation found in a list of 2,674 triliteral roots taken from a dictionary of standard Arabic (Cowan 1979). The statistics in this and all further tables for Arabic were calculated over the same observed counts as used in Frisch *et al*. (2004), based on files supplied by Stefan Frisch (p.c.). These consonant pairs are from adjacent positions in the root, but given Arabic’s root-and-pattern morphological system, they are often separated in the surface phonological string by vowels supplied by other morphemes. ‘Expected’ (E) is the number of such pairs that would occur if the consonants co-occurred at chance level.[1] O/E is the ratio of observed to expected. A consonant pair with an O/E-value below 1 occurs less frequently than expected under a chance combination of consonants, and a pair with an O/E-value over 1 appears more frequently than expected. We will use the terms *overattested* and *underattested* to refer to pairs with O/E values greater and less than one respectively, and will speak of *degree of attestedness* when comparing O/E values of two pairs. Although degree of *representation* is often used in this context (see e.g. Frisch *et al*. 2004), this risks being ambiguous with other uses of the term representation in phonology. In Table 1, we see that the homorganic coronals stand out from the other places of articulation in being less severely underattested.

The degree of attestedness of the Arabic coronals depends on how similar they are on other featural dimensions. Table 2 divides the coronals into three classes: sonorants, obstruent stops, and obstruent fricatives. Pairs of sonorants are highly restricted (O/E = .09), but they co-occur freely with the obstruents (O/E = 1.23 with stops, O/E = 1.20 with fricatives). Pairs of obstruent stops and obstruent fricatives are also highly underattested (O/E = .17 for stops, O/E = .05 for fricatives), but stop-fricative pairs take on an intermediate value (O/E = 0.52).

[1] For each consonantal sequence, ‘expected’ is calculated as the probability of the first consonant as the first member of a sequence, times the probability of the second consonant as the second member of a sequence, times the total number of sequences. For example, for the sequence [t]-[d], the probability of [t] is the total number of sequences with [t] as its first member (102) divided by the total number of sequences (5348; each triliteral root has two sequences), which yields a probability of 0.019. Similarly, the probability of [d] is the total number of sequences with [d] as the second member (243), divided by 5348, yielding 0.045. Expected is then 0.019 * 0.045 * 5348 = 4.63. Like the observed values, the expected values in our tables sum over all sequences with the specified characteristics.

	Coronal Sonorant			Coronal Stop			Coronal Fricative		
	l r n			t d tˤ dˤ			θ ð s z sˤ zˤ ʃ		
Sonorant	12	132.93	**0.09**						
Stop	268	217.21	**1.23**	7	41.81	**0.17**			
Fricative	371	308.36	**1.20**	87	166.01	**0.52**	5	93.66	**0.05**

Table 2: Sequences of non-identical coronals in Arabic verbal roots

McCarthy (1988) accounts for the split between sonorants and obstruents by specifying that the constraint against homorganic segments, the OCP for consonantal place, applies only within the subclasses of coronals defined by the feature [+/-sonorant] (OCP=Obligatory Contour Principle; Leben 1973). In Dresher (1989), Selkirk (1991), and Padgett (1995), the general OCP-PLACE constraint is elaborated into a set of more specific constraints that are violated only by segments that are identical in particular ways (see especially Padgett 1995 on this aspect of the proposal). The relativized OCP constraints that would correspond to McCarthy's (1988) groupings appear in (1).[2] In McCarthy's analysis, the consonants in question are assumed to be adjacent at a derivational level in which intervening vowels are absent.

(1) OCP-LAB — Adjacent labials are prohibited
OCP-DOR — Adjacent dorsals are prohibited
OCP-PHAR — Adjacent pharyngeals are prohibited
OCP-COR[αSON] — Adjacent coronals agreeing in [+/-sonorant] are prohibited

These constraints are categorical: the grammar bans pairs of consonants that contradict them, and permits ones that do not. With OCP-COR relativized to [+/-sonorant], obstruent/sonorant pairs of coronals are permitted, while pairs of obstruents and pairs of sonorants are banned. Under this account, the attested words with sequences of this type would presumably be treated as exceptions.

When constraints apply in this categorical fashion, it is impossible to create a grammar compatible with the intermediate degree of attestedness of the fricative/stop pairs (O/E = .52). Given the constraints in (1), they are ruled out. If the OCP-COR for obstruents is further specified to apply only between segments that have the same value for [+/-continuant], as in Padgett (1995), then there is no restriction at all on stop/fricative pairs. The fact that neither of these seems correct is reflected in McCarthy's (1994) statement that the continuancy restriction on the OCP-COR is "not absolute". McCarthy does not, however, formalize the non-absoluteness of the restriction.

Frisch *et al.* (2004) propose an alternative account of the Arabic place restriction, based on a similarity metric that calculates the similarity of a pair of homorganic segments (x, y) as in (2). The more natural classes two segments share, the closer the value produced by the formula will be to 1, and the more they differ, the closer the value will be to zero. Since, as shown in Table 2, O/E values decrease with increased similarity, O/E and similarity negatively correlate (see section 5.1 on the degree of correlation).

[2] McCarthy (1988) also divides the dorsals by sonority, classifying /w, y/ as the dorsal sonorants, though he states that their lack of co-occurrence may have other explanations.

(2) $$\frac{\text{Number of natural classes shared by } x \text{ and } y}{\text{Number of shared natural classes + number of unshared natural classes}}$$

There are a number of differences between the similarity metric and the OCP-PLACE accounts of Arabic. One is that the similarity metric is gradient, so that it is capable of producing values that correlate with intermediate grades of attestedness. Another is that it calculates similarity based not just on sonority and continuancy, but also in terms of other features, including voice and pharyngealization. Frisch *et al*. (2004) provide evidence, which we discuss in sections 2.3 and 2.4, that similarity on these other dimensions also leads to lower O/E values in homorganic consonantal sequences in Arabic.

In the following section, we show that the Western Austronesian language Muna has a gradient restriction against homorganic consonant sequences that provides strong support for Frisch *et al*.' s claim that [voice] agreement can play a role in determining degree of attestedness. However, the Muna data raise several issues for an analysis in terms of the similarity metric. The first arises from the fact that in displaying such strong evidence for the role of [voice], Muna diverges considerably from Arabic, in which [voice] agreement is a much weaker determinant of co-occurrence rates. We document this divergence by providing a comparison of the role of sonorancy, continuancy, and voice in affecting rates of homorganic consonant co-occurrence in the two languages (sections 2.2, 2.3 and 2.5). The similarity metric provides no mechanism for such cross-linguistic variation. The second problem arises from the fact that Muna, like Arabic, shows greater tolerance for coronal sequences than for labial or dorsal pairs. The similarity metric derives the relative weakness of the Arabic coronal restriction from the size of its coronal inventory. Since there are many more coronals, there are many more natural classes of coronals so that the denominator in the formula in (2) will always be much larger for the coronals than the other places of articulation. This explanation fails to generalize to Muna, which does not have a sufficiently large enough difference between the size of the coronal inventory and that of the other places to yield the differences in O/E values. The place-related differences in the strength of the co-occurrence restriction are documented in section 2.1., and the issues raised for the similarity metric are discussed in section 5.1.

Our own account of gradient restrictions on place co-occurrence is inspired by these earlier analyses, but differs in a number of crucial ways. Following the OCP-PLACE accounts, we posit a set of constraints that are relativized to different places of articulation, and to different subsidiary features. The first difference is that we allow these constraints to be relativized to a larger set of features, including [voice] (the full constraint set is presented in section 4.1). The second difference is that we adopt a theory of constraint interaction that allows constraints to take on a range of values: Harmonic Grammar (Legendre, Miyata and Smolensky 1990/2006, Smolensky and Legendre 2006). We assume a version of Harmonic Grammar that is identical to its better-known sibling Optimality Theory (Prince and Smolensky 1993/2004), except that the constraints are placed along a numerical scale, rather than in a ranking. Input-Output mappings are given a Harmony score based on the sum of their constraint violations, each multiplied by the value, or weight, of the associated constraint. As has been noted in earlier work (e.g. Legendre *et al.* 1990/2006; Keller 2006; Hayes and Wilson 2007), degrees of well-formedness can be expressed in terms of Harmony scores. Section 3.1 surveys experimental research showing that knowledge of phonotactic well-formedness is indeed gradient, and in section 3.2 we introduce our approach and illustrate it with an analysis of the gradient pattern in Arabic coronal co-occurrence we have discussed in this introduction. In section 3.3, we argue for a definition of Harmony-based well-formedness in terms of the *difference* between the score of an Input-Output mapping and the most optimal distinct mapping for the same Input.

In section 4.2, we provide a learnability account of the acquisition of frequency-based weightings of the constraints. Drawing heavily on proposals of Boersma (1998), Boersma and Hayes (2001) and Zuraw (2000), we show that such weightings emerge as the byproduct of gradual learning. We perform learning simulations using Praat (Boersma and Weenink 2007) for both Muna and Arabic, and calculate the correlations between the resulting Harmony scores for homorganic sequences and their O/E values, as well as between the values yielded by Frisch *et al*.'s similarity metric and O/E. For both Muna and Arabic, the Harmony scores yield a better fit to the O/E values. Our analyses of Muna and Arabic are presented in sections 4.4 and 4.5 respectively. Section 5.1 compares these analyses to ones based on the similarity metric, and section 5.2 discusses the relationship of our proposals to other Harmony-based accounts.

2. Consonant co-occurrence patterns in Muna and Arabic

Although place co-occurrence restrictions are best known for their appearance in Semitic languages (Greenberg 1950), there is a growing list of languages from several unrelated families that have been shown to have similar patterns.

(3) a. Javanese (Uhlenbeck 1949; Mester 1986; Yip 1989)

b. English, French, Latin (Berkley 2000)

c. Russian (Padgett 1995)

d. Yamato Japanese (Kawahara *et al.* 2006)

e. Rotuman (McCarthy 2003b)

In this section, we report on the consonant co-occurrence restrictions in the Western Austronesian language Muna (van den Berg 1989). Muna's consonantal inventory is presented in (4).[3]

[3] Muna also has a set of palatals recently borrowed from Indonesian, which are omitted from the inventory in (4), and from all other discussion in this paper. Van den Berg (1989, p. 16) states: "The palatal consonants /c/, /j/, and /y/ are marginal loan phonemes. The number of words containing these recent loan phonemes is very low." He further notes that they are replaced in all but very recent loans with native phonemes.

(4) Consonantal inventory of Muna

	Labial	*Coronal*	*Velar*	*Uvular*	*Glottal*
Voiceless	p	t	k		
Voiced	b	d	g		
Implosive	ɓ	ɗ			
Nasal	m	n	ŋ		
Voiceless prenasal	^{m}p	^{n}t ^{n}s	^{ŋ}k		
Voiced prenasal	^{m}b	^{n}d	^{ŋ}ŋ		
Voiceless fricative	f	s			h
Voiced fricative				ʁ	
Trill		r			
Lateral		l			
Glide	w				

In his grammar of Muna, van den Berg (1989) examines the consonantal co-occurrence patterns in a set of 1100 CVCV roots. He finds restrictions against multiple prenasalized obstruents as in Timugon Murut (Prentice 1971) and in Gurindji (McConvell 1988; Evans 1995), and against pairs of unlike liquids as in Sundanese (Cohn 1992). Van den Berg also notes that homorganic stops that differ in voicing do not co-occur, nor do nasals and homorganic obstruents. These last two gaps are suggestive of a broader generalization: a restriction on non-identical homorganic segments, as observed in Semitic and in the languages in (3).

To further investigate these co-occurrence patterns, we obtained an electronic Muna word list from René van den Berg, based on van den Berg and Sidu's (1996) Muna dictionary. We extracted from this word list all the (V)CVCV and (V)CVCVCV roots, which yielded a total of 5854 roots. Our focus here is on the patterns of co-occurrence found between sequences of consonants in what Frisch *et al.* (2004) call "adjacent" position – that is, ones separated by at most one vowel. The total number of such pairs of consonants in our corpus is 7892; it is greater than the total number of roots, since (V)CVCVCV roots each have two adjacent pairs. We followed Frisch *et al.* (2004) in counting the number of times each sequence occurs in our corpus (Observed), and then in calculating the number of times that each pair would have occurred if consonants combined freely (Expected), and then in determining the ratio between the two (O/E; see footnote 1 on details of the 'Expected' calculation).[4]

In section 2.1, we first show that Muna does indeed have a co-occurrence restriction between homorganic consonants, and that identical consonants are exempt from the restriction. This section also documents the difference between the strength of the restriction for dorsals, labials and coronals. Section 2.2 then shows how agreement in sonorancy, voicing, and stricture contribute to the strength of the co-occurrence restrictions. In section 2.3, we compare the

[4] The observed and expected values for all the pairs of consonants are available on request from the authors, or at www.umich.edu/~coetzee/Muna_Arabic/.

contribution of these ‘subsidiary’ features in Muna and Arabic. In section 2.4, we discuss the role of further features in contributing to the underattestedness of homorganic consonants in Arabic and Muna. Section 2.5 presents an analysis of the contribution of the entire set of subsidiary features to degree of attestedness of homorganic consonants in both Arabic and Muna, which uses the statistical technique of linear regression.

2.1. Homorganicity and identity in Muna

Restrictions on place feature co-occurrence, as well as those on laryngeal feature co-occurrence (MacEachern 1999), sometimes fail to apply to identical consonants. We start by showing that identical consonants in Muna co-occur freely, before turning to the restriction on non-identical homorganic consonants.

Muna has 18 pairs of identical consonants,[5] and 15 of these have O/E values above 1, showing that pairs of identical consonants typically occur more frequently than expected by chance. We tested the statistical significance of this tendency with a *t*-test, and found that the O/E-values for pairs of identical consonants are significantly higher than 1 ($t(17) = 4.10$, $p < .001$, one-tailed, $\mu = 2.01$, $\sigma = 1.05$). For this reason, we exclude pairs of identical consonants from all further tables and calculations related to the place co-occurrence restrictions.

Although the identity exemption in Muna is familiar from other languages, there is a way in which the Muna case is special. Consonants that are identical except that one is pre-nasalized (e.g. [k-$^{\text{ŋ}}$k]) might be taken to differ in only one feature, and to be therefore highly similar in terms of their featural makeup. The expectation would then be that they should be highly restricted in their co-occurrence. They can be compared with homorganic pairs that differ in terms of prenasalization and one additional feature such as voicing (e.g. [k-$^{\text{ŋ}}$g]). Such pairs differ in two features and are hence less similar, so that the expectation is that they would be less severely restricted and occur more freely than pairs that differ only in terms of prenasalization. This expectation is not borne out. In fact, the opposite is observed – pairs of consonants that differ only in prenasalization occur more freely than pairs that differ in prenasalization and voicing. There is a total of six such pairs of pairs of consonants: [^{m}p~p] vs. [^{m}p~b], [^{m}b~b] vs. [^{m}b~p], [^{n}t~t] vs. [^{n}t~d], [^{n}d~d] vs. [^{n}d~t], [$^{\text{ŋ}}$k~k] vs. [$^{\text{ŋ}}$k~g], [$^{\text{ŋ}}$g~g] vs. [$^{\text{ŋ}}$g~k]. In 5 out the 6 instances, the O/E-value for the consonants that differ only in prenasalization (e.g. [k-$^{\text{ŋ}}$k]) is higher than that of consonants that differ in both pre-nasalization and voicing (e.g. [g-$^{\text{ŋ}}$k]).[6] The difference between these two kinds of consonant pairs reaches significance under a Wilcoxon Rank-Sum test: $Z = 1.78$, $n = 6$, $p < .04$ (Mann and Whitney 1947). Thus, pairs of homorganic consonants that differ only in prenasalization seem to pattern more like identical consonants than like non-identical consonants, perhaps because identity is computed on the oral portion of the complex segment. Like the pairs of identical consonants, we exclude these pairs from further tables and calculations.

---

[5] We exclude from this, and from all further tables and calculations pairs of prenasalized consonants, which as van den Berg (1989) shows, are subject to an independent absolute restriction.

[6] As discussed above for Arabic, throughout the paper, we abstract from the order of the consonants in a pair when we calculate the O/E values of non-identical pairs of consonants. Following Frisch *et al.* (2004) we sum observed and expected values for the two sequences involving a pair, and then calculate O/E from the summed values.

Further study is required to better understand the nature of the identity exemption. In particular, in many cases the identical consonants do precede identical vowels, suggesting some form of reduplication. We did exclude all affixes that were marked as such in van den Berg and Sidu (1996) from our corpus; nonetheless, it is still possible that some of the identical pairs are diachronically due to reduplication (see Yip 1989 on Javanese), or that Muna is subject to what Zuraw (2002) calls 'Aggressive Reduplication'.

Non-identical homorganic consonants are restricted in their co-occurrence; this becomes clear in a comparison with non-identical heterorganic consonants, which co-occur relatively freely. Table 3 presents the O/E-values for non-identical consonants by place of articulation. The observed and expected values were each summed for all sequences that fall into each category, and then the O/E ratios were computed from the sums. Inspection of the table will confirm that the values within place classes are all below 1, while the values across places of articulation are all above 1.

	Labial			**Coronal**			**Dorsal**		
	Obs	Exp	**O/E**	Obs	Exp	**O/E**	Obs	Exp	**O/E**
Labial	132	441.9	**0.30**						
Coronal	2741	2154.37	**1.27**	1338	1686.0	**0.79**			
Dorsal	875	770.9	**1.14**	1780	1527.6	**1.17**	29	163.6	**0.18**

Table 3: Adjacent homorganic and heterorganic consonants in Muna

In our corpus, 77% of pairs of homorganic consonants have O/E-values below 1, compared to 36% of pairs of heterorganic consonants. The fact that homorganic consonants but not heterorganic consonants are restricted in co-occurrence is also confirmed by *t*-tests. Taking the O/E values of each of the pairs of consonants as data points, we found that the O/E-values for pairs of heterorganic consonants are distributed significantly above 1 ($t(176) = 4.7$, $p < .001$, one-tailed, $\mu = 1.2$, $\sigma = 0.55$), while those for pairs of homorganic consonants are significantly below 1 ($t(69) = 9.8$, $p < .001$, one-tailed, $\mu = 0.56$, $\sigma = 0.48$). Since the coronal pairs are the closest to 1 in their overall O/E value, we also ran a *t*-test to determine whether their values are distributed significantly below 1, and found that they indeed are ($t(38) = 2.58$, $p < .008$, one-tailed, $\mu = 0.81$, $\sigma = 0.47$).

As Table 3 shows, pairs of homorganic consonants can be ordered as follows in terms of their O/E-values: coronal > labial > dorsal. It therefore seems to be the case that coronals are least restricted in their co-occurrence, with labials in an intermediate position and dorsals the most restricted. In order to test the statistical significance of this trend, we collected pairs of non-identical consonant pairs that differ only in place of articulation (for example, the voiced-voiceless stop pairs [p-b] and [t-d] differ only in labial vs. coronal place). This controls for the influence of other featural differences on co-occurrence rates. These pairs include the pairs of voiced and voiceless stops ([p-b], [t-d], [k-g]), pairs of voiced stops and nasals ([b-m], [d-n], [g-ŋ]), pairs of voiceless stops and nasals ([p-m], [t-n], [k-ŋ]), pairs of voiced stops and voiceless prenasalised stops ([b-^{m}p], [d-^{n}t], [k-^{ŋ}g]), and pairs of voiceless stops and voiced prenasalized stops ([p-^{m}b], [t-^{n}d], [k-^{ŋ}g]). In the comparison of the labials and coronals, we also included all combinations of the fricative with other consonants ([f-p, f-b, f-m, f-^{m}p, f-^{m}b] and [s-t, s-d, s-n, s-^{n}t, s-^{n}d]). Since the dorsal fricative [ʁ] is voiced, it cannot be compared with the labial and coronal fricatives. We then submitted the pairs of consonant pairs that can be compared between coronal and labial, coronal and dorsal, and dorsal and labial place to the non-parametric Wilcoxon Rank-Sum test in order to test whether the differences in O/E-values between the places of articulation are statistically significant. We opted for a non-parametric test rather than a *t*-test here, since *n* is very low in these tests, and since the data are not normally distributed (especially for the dorsals and labials, there are several consonant pairs with O/E-values of zero). All three comparisons are significant under a one-tailed Wilcoxon: (i) labial vs. coronal, $Z = 2.38$, $n = 10$, $p < .01$; (ii) labial vs. dorsal, $Z = 2.02$, $n = 5$, $p < .03$; (ii) coronal vs. dorsal, $Z = 2.02$, $n = 5$, $p < .03$. This means that the O/E-values for pairs of coronals are indeed higher than that for pairs of labials, which are in turn higher than that for pairs of dorsals.

Although our analysis focuses on the consonants in "adjacent" position, we also examined the co-occurrence rates for consonants in first and third position of a tri-consonantal root, which are separated by an intervening consonant. There was a total of 2038 such pairs in our corpus. As has been reported for Arabic, in Muna there is a restriction on homorganic dorsals and labials in these positions that is somewhat weaker than the one that applies between adjacent consonants (see esp. McCarthy 1988, Pierrehumbert 1993). Non-identical dorsal pairs in these positions have an overall O/E of 0.69 (O=37, E=54), and non-identical labial sequences have overall O/E of 0.71 (O=102, E=143). In both cases the O/E values are below 1, but above the values in adjacent position (0.18 for dorsals, and 0.30 for labials). Non-identical coronals, however, are slightly overattested (O/E=1.12; O=233, E=208). Taking the O/E values of every non-identical dorsal and labial pair as data points, *t*-tests for both places indicates that they are distributed significantly below 1 (Dorsals $t(11) = 2.29$, $p < .05$, one-tailed, $\mu = 0.56$, $\sigma = 0.48$; Labials $t(24) = 5.34$, $p < .01$, one-tailed, $\mu = 0.62$, $\sigma = 0.58$). Contrary to Arabic, identical consonants do not seem to be underattested in these positions in Muna. Their overall O/E is 0.78 (O=95, E=121), which is below 1. However, this result is highly influenced by the low observed values of two coronal pairs that have high expected values: [l-l] and [s-s]. The mean of the O/E values for the identical pairs is in fact slightly above 1 ($\mu=1.13$, $\sigma = 1.24$), and the overall O/E value for the identical dorsals is 0.95 (O=48, E=51), and for the identical labials is 1.02 (O=25, E=25). Thus, the results from the non-adjacent positions confirm the general pattern for adjacent positions: non-identical dorsals and labials are significantly underattested, and identical consonants are not subject to the restriction.

2.2. The influence of sonorancy, voicing, and stricture on place co-occurence in Muna

Having established that Muna has a co-occurrence restriction against non-identical homorganic consonants, and that the strength of the restriction differs across place of articulation, we now turn to the influence of other features. In this section, we show that homorganic pairs of consonants that agree in terms of sonorancy, voice or stricture occur less freely than pairs that do not share these features.

2.2.1 Exemplification: minimally different sets of sequences

Because Muna differs from Arabic in that it does permit labial and dorsal sequences (they are all at or near zero in Arabic), the influence of the subsidiary features on place co-occurrence can be seen across the places of articulation. This is exemplified in the Table 4, which compares the O/E values for homorganic fricatives, stops and nasals at each place of articulation. For the existing pairs, we include examples and glosses. To aid legibility, we leave out the observed and expected values.

		Coronal	Labial	Dorsal
a.	**Nasal + Voiced Stop**	n-d: 0.25 da:no 'true'	m-b: 0.07 bomu 'bomb'	ŋ-g: 0.00 –
b.	**Voiced Stop + Voiceless Stop**	d-t: 0.60 datu 'cross-eyed'	b-p: 0.10 pabu 'ineffective'	g-k: 0.07 kagala 'anklet'
c.	**Nasal + Voiceless Stop**	n-t: 0.70 tuna 'twig'	m-p: 0.39 mopi 'loser in game of Jacks'	ŋ-k: 0.10 kaŋia 'starfish'
d.	**Fricative + Voiceless Stop**	s-t: 0.37 tisore 'run aground'	f-p: 0.22 fopanto 'throw down'	ʁ-k: 0.40 kaʁa 'choking sound'
e.	**Fricative + Voiced Stop**	s-d: 0.55 sida 'happen'	f-b: 0.58 febuni 'hide oneself'	ʁ-g: 0.00 –
f.	**Nasal + Fricative**	n-s: 1.17 nasa 'tired'	m-f: 1.04 mafaka 'agreement'	ŋ-ʁ: 0.00 –

Table 4: O/E-values for Muna voiced and voiceless stops, fricatives, and nasals

We can first note that this chart provides sets of minimally different sequences that illustrate the coronal > labial > dorsal ordering discussed in the last section. In almost every row, the O/E value for the coronals are greater than that of the labials, which are themselves greater than that of the dorsals.

For the subsidiary features, particularly striking is the role of [voice] agreement in affecting relative O/E-values. This can be seen clearly by comparing the values amongst the rows (a, b, c), and amongst (d, e, f). The nasal-voiced stop pairs in row (a), which contain two [+voice] consonants, are more underrepresented than either the voiced-voiceless stop pairs (row (b)), or the nasal-voiceless stop pairs (row (c)). The fact that the nasal-voiced stop pairs, which agree in voicing but not sonorancy, have lower O/E values than the voiced stop-voiceless stops pairs, which agree in sonorancy, may be particularly surprising in light of what is observed in Arabic, in which agreement in sonorancy is a much stronger determinant of underattestedness than agreement in voicing.

Comparing across rows (d), (e) and (f) also reveals voice-related differences. The voiceless fricatives [f] and [s] co-occur less freely with the voiceless stops (row (d)) than with the voiced stops (row (e)) or with the nasals (row (f)). On the other hand, the voiced fricative [ʁ] never occurs with the voiced stop (row (e)) or with the nasal (row (f)), but co-occurs at an intermediate rate with the voiceless stop (row (d)).

It is worth noting that voicing is non-contrastive on both the nasals and the fricatives. The Muna data therefore seems to be problematic for a theory with lexical underspecification of non-contrastive voicing (see e.g. Ito and Mester 1986), since the place co-occurrence restriction is lexical, and non-contrastive voicing plays a role in determining relative rates of underattestedness (see also Frisch *et al*. 2004).

Table 4 also provides some evidence for the role of continuancy agreement: the nasal-voiceless stop pairs are underattested (row (c)), while the nasal-voiceless fricative pairs have O/E-values above 1 (row (f)). Sonority seems to have an effect as well: the voiced stop-voiceless fricative pairs (row (e)) are underattested though the nasal-voiceless fricative pairs are not (row (f)), and the oral stop pairs (row (b)) occur somewhat less freely than the nasal-voiceless stop pairs (row (c)).

2.2.2 Descriptive statistics for the entire dataset

To examine the effect of these features more systematically, we calculated the summed observed and expected values for homorganic non-identical pairs that agree for these features, and for those that do not, as well as the resulting O/E value. The 'Agree-[voice]' category includes pairs of voiced obstruents, pairs of voiceless obstruents, pairs of sonorants, and nasal-voiced obstruent pairs. Voicing of prenasals was determined by the oral portion, and the implosives were taken to be neither [+voice] nor [-voice]. The set 'Others' thus includes sequences with opposite values for [+/-voice], as well as sequences in which one of the members has no [voice] value. The 'Agree-[sonorant]' category includes sonorant pairs and obstruent pairs, with sonorancy again determined by the oral portion of the prenasals. For stricture, we sought an alternative to the feature [continuant] (cf. Frisch *et al.* 2004), because there is no evidence in either Arabic or Muna that fricatives co-occur less freely with homorganic liquids than do stops. Such an effect would be expected if they were included in a single [+continuant] class for the place co-occurrence restriction. We therefore created a category for consonant pairs that have the same degree of oral stricture, which we call 'Agree-[stricture]'. This category includes pairs of stops (including both nasal and oral consonants), pairs of fricatives, and pairs of liquids. The stricture of the prenasalized fricative was determined by the fricative. However, because we treat [^{n}s] and [s] are identical for the place co-occurence restriction, and because there are no other homorganic fricative pairs, no fricatives are included in the Agree-[stricture] category. Table 5 contains the results for Muna.

	Agree			**Others**		
	O	**E**	**O/E**	**O**	**E**	**O/E**
[sonorant]	381	973.18	0.39	1118	1318.40	0.85
[voice]	361	761.54	0.47	1138	1530.04	0.74
[stricture]	245	713.63	0.34	1254	1577.95	0.79

Table 5: Muna non-identical homorganic consonants with subsidiary feature agreement

In all cases, the O/E-values for the pairs in the 'Agree' category are substantially lower than those in the 'Others', thus supporting the observations made for Table 4 about the role of these features in lowering O/E-values. However, the fact that these categories overlap makes it difficult to assess the independent contribution each of these features to degree of attestedness. Table 6 provides a measure of the overlap by showing the percentage of the pairs in each category that also belong to another. The percentages refer to the proportion of the members of the category indicated in the topmost row that belong to the one indicated in the second row.

Agree-[sonorant]		**Agree-[voice]**		**Agree-[stricture]**	
Agree-[voice]	**Agree-[stric]**	**Agree-[son]**	**Agree-[stric]**	**Agree-[son]**	**Agree-[voice]**
31%	50%	50%	30%	65%	24%

Table 6: Percentage of pairs agreeing in [sonorant], [voice] or [stricture] that also agree in one of the other two features

Because of the high degree of overlap, one cannot conclude from a comparison of the O/E values in one of the 'Agree' categories to 'Other' in Table 5 that the feature makes an independent contribution (though see section 2.5 for a regression analysis that suggests that all in fact do). The numbers in Table 6 also show that it would be a mistake to conclude from Table 5 that [voice] agreement has a weaker effect than [stricture] or [sonority] on O/E-values. A higher percentage of the members of Agree-[sonorant] than Agree-[voice] also belong to Agree-[stricture], and a higher percentage of the members of Agree-[stricture] than Agree-[voice] also belong to Agree-[sonorant]. Thus, the apparently stronger effect of agreement for [sonorant] and [stricture] might instead be due to a cumulative effect of the two features. This is especially likely given that the data contained in Table 4 point to such a strong influence of [voice].

To get a more precise descriptive picture of the influence of each of the features, we subdivided each of the categories in terms of whether the pairs also agreed for the other features. The resulting co-occurrence statistics are presented in Table 7. We continue to use the label 'Others' for the 'Non-Agree-Voice' category, since it contains the unspecified implosives. For the other features, all segments are specified for a value, so we use the label 'Differ'.

	Agree-Sonorant						**Differ-Sonorant**					
	Agree-Stricture			**Differ-Stricture**			**Agree-Stricture**			**Differ-Stricture**		
Agree-Voice	30	155.21	**0.19**	107	301.90	**0.35**	19	77.03	**0.25**	205	227.40	**0.90**
Others	136	317.18	**0.43**	108	198.89	**0.54**	60	164.22	**0.37**	834	849.75	**0.98**

Table 7: Muna non-identical homorganic pairs by subsidiary feature category (O, E, O/E)

This table indicates independent and cumulative effects for each of the features. The lowest degree of attestedness is observed for the segment pairs that agree for all three features (there is in fact only one such pair, [l-r]), and the highest degree of attestedness is for those that agree in none of them. Every pairwise comparison between between corresponding "Agree" and "Others" cells for voicing, show a lower degree of attestedness for the "Agree" than the "Others" cell. Similarly, every pairwise comparison between corresponding "Agree-Stricture" and "Differ-Stricture" cells show lower attestedness in the "Agree" than the "Differ" cell. With one exception this is also true for sonorant. The one exception is in the comparison between "Agree-Sononant, Agree-Stricture, Others-Voice" and "Differ-Sononant, Agree-Stricture, Others-Voice", where the "Differ-Sonorant" cell has a slightly lower O/E-value (0.37) than the "Agree-Sonorant" cell (0.43).

In sum, Muna place co-occurrence appears to be affected by all three of the subsidiary features [sonorant], [voice], and [stricture]. No one feature seems to dominate. This contrasts with Arabic, as well as the other languages listed in (3), in which [sonorant] agreement has always been observed to have the strongest effect on co-occurrence. In the next section, we reproduce these descriptive analyses using data from Arabic, in order to illustrate the differences between the two languages.

Before turning to Arabic, though, we make one more observation about the Muna data. Since Muna is unique in the extent to which disagreement in [voice] functions to weaken the place co-occurrence restriction, one might seek an alternative explanation for the lower O/E-values seen in pairs that agree in [voice]. One such alternative could be that [voice] disagreement is general in the language, and not specific to homorganic pairs. This does not seem to be the case, however. Heterorganic pairs that agree in [voice] are overattested (O = 2310, E = 1994.57, O/E = 1.16), though they are slightly less overattested than pairs that differ in [voice] (O = 2395, E = 1902.91, O/E = 1.26).

2.3. The influence of sonorancy, voicing, and stricture on place co-occurence in Arabic

The tables in this section classify the Arabic coronal pairs using the same categories as for the Muna homorganic pairs. To make the comparison as close as possible, we leave out the identical pairs, and consider only pairs from adjacent positions in the root. Like the data discussed in the introduction, the data here come from Cowan (1979) via Stefan Frisch (p.c.), and are thus based on the same corpus used by Frisch *et al.* (2004). Table 8 presents the summed observed and expected values for the sequences in each category, along with the O/E values derived from them (in bold). Since all segments in Arabic are specified as [+/-voice], we use the category label 'Differ-Voice' here.

	Agree-Sonorant						**Differ-Sonorant**					
	Agree-Stricture			**Differ-Stricture**			**Agree-Stricture**			**Differ-Stricture**		
Agree-Voice	3	111.82	**0.03**	49	160.36	**0.31**	37	31.71	**1.17**	202	174.77	**1.16**
Differ-Voice	9	76.04	**0.12**	50	86.19	**0.58**	34	26.46	**1.28**	386	308.66	**1.25**

Table 8: Arabic non-identical coronal pairs by subsidiary feature category (O, E, O/E)

As in Muna, the O/E-values in the 'Agree' categories are consistently lower than their 'Differ' counterparts, showing a role for each of the features. The difference between Muna and Arabic is that Arabic shows a clear hierarchy of subsidiary feature strength. Whether or not a pair agrees in [sonorant] is undoubtedly the strongest determinant: the four ‘Agree-Sonorant’ categories have O/E values indicating that they are underattested, while the ‘Differ-Sonorant’ categories all have values indicating that they are overattested. The primary role of sonorancy and the secondary role of obstruent continuancy in influencing Arabic coronal co-occurrence have long been recognized (see the discussion in section 1). In what follows, we look more closely at the relative strength of continuancy and voicing agreement.

Frisch *et al.* (2004) discuss the role of [voice] agreement in affecting co-occurence rates both within the obstruents and across the sonorancy divide. For the obstruents, they note a difference in O/E rates for homorganic pairs (both adjacent and non-adjacent) that agree in voicing (0.21), versus those that do not (0.36), but do not report a statistical test of this difference (p. 192). For the obstruent-sonorant pairs, they find differences in O/E values corresponding to those between the ‘Agree-Voice’ and ‘Differ-Voice’ rows under the ‘Differ-Sonorant’ column in Table 8. They also performed a Wilcoxon test on obstruent-sonorant pairs that differed only in whether they agree for voice (e.g. [t-l] vs. [d-l]), which nearly reaches significance ($p < 0.07$). They conclude that the effect of voicing in Arabic place co-occurrence is “small” and “subtle” (p. 195).

For the sonorant-obstruent pairs, it is difficult to compare the relative effect of voicing and stricture agreement, since there are only 4 pairs in the ‘Agree-Stricture’ category (the nasal-stop pairs). In the ‘Agree-Sonorant’ categories, there are more data to draw on. Some of the difference between ‘Agree-Voice’ and ‘Differ-Voice’ in Table 8 is due to the fact that ‘Agree-Voice’ contains [+sonorant] pairs, which are high in Expected and low in Observed values, but ‘Differ-Voice’ does not. Revised figures that omit the sonorants appear in Table 9. Within the obstruents, we continue to see a difference between the pairs that agree in [voice] and those that do not, though the difference is much less marked than that caused by stricture agreement.

	Agree-Stricture			**Differ-Stricture**		
Agree-Voice	3	59.43	**0.05**	37	79.82	**0.46**
Differ-Voice	9	76.04	**0.12**	50	86.19	**0.58**

Table 9: Arabic non-identical coronal obstruent pairs

In sum, Arabic, like Muna, provides evidence that agreement for sonorancy, voicing, and stricture increases the strength of the restriction against homorganic pairs. Arabic differs, though, in that [sonorant] is clearly the strongest determinant of degree of attestedness, and [voice] is clearly the weakest.

2.4 Further subsidiary features

There are further features that are specific to each of Muna and Arabic that appear to affect co-occurrence rates of homorganic segments. In this section we document their role, before turning to a statistical analysis of the contribution of the entire set of subsidiary features in each language.

In Muna, there is evidence that agreement in prenasalization lowers O/E values. We showed above (section 2.1) that pairs that are identical except that the one is prenasalized (e.g. [k-^{ŋ}k]) occur more freely than pairs that differ both in prenasalization and voicing (e.g. [g-^{ŋ}k]). This leads us to conclude that pairs like [k-^{ŋ}k] are treated as identical pairs in spite of their difference in prenasalization. However, this does not mean that prenasalization is completely inert in determining the patterns of consonant co-occurrence in Muna. When a prenasalized segment is paired with a homorganic segment that is not identical to its oral portion (e.g. [k-^{ŋ}g]), the difference in prenasalization does seem to mitigate the strength of the co-occurrence restriction – these pairs are better attested than the corresponding pairs that agree in prenasalization (e.g. [k-g]). The values for all these pairs that can be compared across place of articulation are shown in Table 10. This same effect is seen in how the coronal fricatives [s] and [^{n}s] co-occur with [t] and [d]. The pair [d-s] that agrees in prenasalization has a lower O/E-value than the pair [d-^{n}s] that differs in prenasalization (0.55 vs. 1.68). However, agreement in [voice] trumps agreement in prenasalization: the pair [t-^{n}s] has no advantage over the pair [t-s] in spite of the fact that [t-^{n}s] differs in prenasalization. In fact, [t-^{n}s] is less well attested than [t-s] (0.05 vs. 0.37). The fact that [t-^{n}s] agrees in voicing seems to cancel out any advantage that this pair has because of its difference in prenasalization.

Agree in prenasalization?	Pair	Observed	Expected	O/E
Yes	p~b	2	20.01	0.10
No	^{m}p~b, ^{m}b~p	12	41.74	0.29
Yes	t~d	26	43.57	0.60
No	^{n}t~d, ^{n}d~t	54	48.21	1.12
Yes	k~g	2	30.33	0.07
No	^{ŋ}k~g, ^{ŋ}g~k	2	19.41	0.10

Table 10: Pairs that differ in voicing, and that either agree or differ in prenasalization.

In Arabic, Frisch *et al*. (2004: 192) point out that the emphatic coronals never co-occur with each other in adjacent position. Table 11 shows that pairs that differ in the feature [pharyngeal] and pairs that both specified as [-pharyngeal] are observed, while pairs of [+pharyngeal] consonants are not.

	[+phar]			[-phar]		
	O	E	O/E	O	E	O/E
[+phar]	0	18.16	**0**			
[-phar]	48	126.63	**0.38**	51	156.69	**0.33**

Table 11: Co-occurrence of Arabic coronal obstruents in terms the feature [pharyngeal]

2.5 Statistical analysis of the contribution of the subsidiary features

As discussed in section 2.2, the overlap between classes defined by the subsidiary features makes it difficult to determine whether each one makes an independent contribution to lowering O/E values. In this section, we present an analysis that tests the contribution of each feature while taking the others into account, using the statistical technique of multiple linear regression. In this we follow Kawahara (2007) who analyze the contribution of multiple features on the co-occurrence of consonants in Japanese using a similar technique.[7]

A linear regression aims to find the coefficient values that best fit a set of equations. One side of the equations contains the values of the dependent or response variable. The dependent variable in our analysis was a transformed Observed value for each of the consonant pairs: its natural logarithm (the usual frequency statistic in studies of lexical frequency). Because a log-value cannot be computed for zero, we added 0.5 to all values before taking the logarithm. The other side of the equations contains the explanatory variables, along with the coefficients that the regression adjusts. One of the explanatory variables in our analysis was the natural logarithm of the Expected value. If no other factors besides the individual segmental frequencies affected the co-occurrence rates of the consonant pairs, then Observed should match Expected (modulo random factors). The other explanatory variables were the 'Agree' classifications discussed in the previous sections: if a pair of segments shared the feature specification, they received a value of '1', else '0'.[8] The model for a sequence *xy* is shown in (5), where 'β_n' is a coefficient, and the abbreviated label for each of the categories stands for the value of the sequence on that variable. Based on the observations in the previous sections, we expect that pairs that have a value of '1' for the 'Agree' variables should generally have lower values for Observed than Expected (i.e. O/E values beneath 1). Therefore, the coefficients of these variables should be negative, to bring the values of the right-hand side of the equations closer to those of the left-hand side.

---

[7] The particular regression model we use is somewhat different from that of Kawahara (2007); thanks to Adam Albright and Hugo Quené for useful suggestions in its development.

[8] Since the independent variables in this regression are a mixture of continuous and categorical variables, the analysis performed here is an ANCOVA (analysis of covariance) rather than an ordinary linear regression. We performed this analysis in SPSS 14.0, which automatically performs an ANCOVA when the independent variables are of a mixed kind.

(5) *Regression model*

$$\log.O_{xy} = \beta_1 \log.E_{xy} + \beta_2 \text{AGR-SON}_{xy} + \beta_3 \text{AGR-VCE}_{xy} + \beta_4 \text{AGR-STRIC}_{xy}$$
$$+ (\beta_5 \text{AGR-PRENAS}_{xy}) \textit{ Muna only}$$
$$+(\beta_5 \text{AGR-EMPH}_{xy}) \textit{ Arabic only}$$

For Muna, all of the non-identical homorganic pairs were included in the analysis. The output of the regression analysis for Muna is shown in Table 12. The R^2 for this regression is 0.75 ($F(5, 267) = 156.0$, $p < .001$). In this table, we label the Agree variables with the feature they assess. As expected, the coefficients for all of these variables are negative. Furthermore, they all have relatively similar values.[9]

	Coefficients	*Standard Error*	*t Stat*	*p-value*	*Tolerance*
Expected	1.15	0.05	24.51	< .01	0.95
Sonorant	-0.22	0.06	-3.93	< .01	0.74
Voice	-0.25	0.07	-3.76	< .01	0.84
Stricture	-0.31	0.06	-5.02	< .01	0.79
Prenasal	-0.28	0.06	-4.87	< .01	0.70

Table 12: Linear regression for Muna non-identical homorganic consonants

Table 12 also shows the results for significance tests that examine the contribution of each of the other features when all of the others are included (that is, the difference between the fit of the equations when the variable is included, and when it is left out). All of the variables reach significance in this analysis as indicated by *p* values beneath 0.01.

The final column in Table 12 gives the tolerance value for each of the coefficients. The tolerance of an independent variable var_1 is expressed as $(1 - R^2)$, where R^2 is the result of regressing var_1 on the remaining independent variables. If the tolerance for some independent variable is 0 (i.e. $R^2 = 1$), it means that that independent variable is a perfect linear combination of the other independent variables already in the regression equation – it is therefore perfectly collinear with the other variables. As shown in Table 12, all of the predictor variables have relatively high tolerance values, implying that collinearity is at an acceptable level.[10]

The Arabic results in Table 13 look quite different from that of Muna. Here, we included just the non-identical coronal pairs, since homorganic pairs at the other places are of too low frequency to show the effects of the subsidiary features. The R^2 value is 0.79 ($F(5, 85) = 57.6$, $p < .001$).

---

9 The files that we used to perform these regressions for both Muna and Arabic are available at www.umich.edu/ ~coetzee/Muna_Arabic/.

10 It is usually accepted that tolerance values below 0.1 indicate unacceptably high collinearity (Faraway 2005:83-87; Miles 2005). Since all of the tolerance values in this regression are equal to or higher than 0.7, we conclude that the collinearity is not too high in this analysis.

	Coefficients	*Standard Error*	*t Stat*	*p-value*	*Tolerance*
Expected	0.63	0.11	6.01	< .01	0.68
Voice	-0.09	0.08	-1.23	.22	0.99
Sonorant	-0.87	0.10	-9.15	< .01	0.68
Stricture	-0.17	0.09	-2.01	< .05	0.85
Emphatic	-0.03	0.08	-.36	.72	0.91

Table 13: Linear regression for Arabic coronals

In comparison with Muna, sonorant stands out much more dramatically from the other features. Its coefficient is considerably lower, and it is the only one of the features that reaches significance at the 0.01 level (the Expected value is not a 'feature' in the relevant sense). Also note that voicing does not reach significance at all, showing that agreement in voicing alone does not contribute significantly to the co-occurrence restrictions. Although there is overlap in the predictor variables as discussed earlier, the tolerance levels for all of the predictors are high, so that collinearity is at an acceptable level.

2.6 Conclusions of descriptive analysis

In this section, we have shown that Muna, like Arabic (and many other languages) has a pattern of underattestedness of homorganic sequences of consonants. As in Arabic, similarity on other featural dimensions contributes to the degree of attestedness of homorganic sequences. However, the relative contribution of these subsidiary features differs in the two languages: Arabic shows an overwhelming effect of sonorancy agreement in lowering attestedness, while Muna has a more balanced contribution of voicing, sonorancy and stricture.

Another way in which Muna resembles Arabic is in the gradience of the place co-occurence restriction. As Frisch *et al.* (2004) point out, standard generative phonology, including OT, usually provides no mechanism for capturing this sort of gradience in the grammar. Therefore, before turning to our phonological analysis of these patterns, we first discuss general theoretical issues surrounding the treatment of gradient well-formedness in Optimality Theory and Harmonic Grammar.

3. Gradient acceptability in Harmonic Grammar

In the version of OT proposed by Prince and Smolensky (1993/2004), the output of the grammar is categorical. The optimal form is grammatical, or well-formed, and sub-optimal candidates are ungrammatical, or ill-formed (p. 4). There is no distinction either between degrees of well-formedness, or between degrees of ill-formedness. In this section, we begin by briefly reviewing experimental work that provides evidence that speakers do make finer distinctions amongst consonantal sequences: that some sequences that are attested in a language are preferred over other attested ones, and that some unattested sequences are preferred over other unattested ones. We then discuss how the numerically weighted constraints in Harmonic Grammar can be used to model gradient phonotactics, and compare this approach with some proposed extensions of OT that aim to deal with gradient acceptability.

3.1 Gradient phonotactic acceptability

The best-studied case of gradient phonotactic acceptability is that found amongst English onset clusters. Within the set of clusters that are unattested in English words, there are relatively clear distinctions in their degrees of ill-formedness. These distinctions have been found in judgment studies (e.g. Scholes 1966, Pertz and Bever 1975, Albright 2007ab, Hammond 2007), in studies of production (e.g. Davidson 2006) and in studies of perception (e.g. Berent *et al*. 2007). For example, the Berent *et al*. study shows that ill-formed onset clusters with rising sonority (e.g. [pn]) are more accurately perceived than clusters with flat sonority (e.g. [pt]), which are themselves perceived more accurately than clusters with falling sonority (e.g. [lp]). If we follow Berent *et al*. in interpreting this experiment as tapping phonological knowledge (cf. Peperkamp 2007), then we have evidence for three degrees of ill-formedness: [lp] > [pt] > [pn].

English onsets also provide examples of gradient well-formedness, that is, distinctions between existing structures. For example, standard American English does seem to tolerate labial stop-glide clusters – they are retained in borrowings like *pueblo*, *Puerto Rico*, and *bueno.* Hayes and Wilson (2007) label this sequence as 'exotic', and it is clear that it has a marginal status compared to well-attested clusters like [pl]. It is not, however, of the same status as completely unattested clusters (see Moreton 2002). Thus, we can add two further degrees to the English onset ill-formedness scale: [lp] > [pt] > [pn] > [pw] > [pl].

As the Muna and Arabic place restrictions are instances of gradient attestedness, it is important to note that there have been several experimental studies that show that speakers do distinguish between the acceptability of existing sequence types in their language. Berent and Shimron (1997) and Berent *et al*. (2001) asked Hebrew speakers for word-likeness judgments of nonce words that contained three consonants. Their stimuli contained one group of words with identical consonants in the first two positions of the consonantal root (SSM). Like Arabic, Hebrew does not allow roots of this type, so this group of nonce words represents a non-occurring root type in Hebrew. But their stimuli also included two groups of roots with only consonant sequences that actually occur in Hebrew – one group with no identical consonants (PSM), and another with identical consonants in the final two root positions (SMM). Their subjects rated the nonce words with non-occurring sequences (SSM) as less well-formed than both of the groups with only occurring sequences. Importantly, they also distinguished between the two groups with occurring sequences – they rated the PSM-type nonce word better than the SMM-type. Other experimental data on gradient well-formedness include Frisch and Zawaydeh (2001) on Arabic, and Bailey and Hahn (2001), Coetzee (2004, 2008, to appear), Greenberg and Jenkins (1964), Hammond (2004), Treiman *et al*. (2000) and others on English.

Most work on phonotactics in generative phonology, in OT and elsewhere, assumes a two-way distinction between ill-formed and well-formed sequences (see recently Hayes 2004 and Prince and Tesar 2004). As far as we know, no argument has ever been offered that this limitation has any principled basis. Even though the two-way grammaticality distinction seems to have guided practice in generative phonology, Chomsky and Halle in fact propose an account of gradient phonotactics in SPE (1968: 417; see Ohala and Ohala 1986 for a critique). The experimental research we have just reviewed, and much more, suggests that speakers do make distinctions amongst phonological forms that go beyond a simple grammatical/ungrammatical split. These distinctions also seem to go beyond the three-way categorization that would be offered by including 'exceptional' as an intermediate category, as is sometimes done in standard generative phonology. In the next section, we show how gradient phonotactics can be captured in a version of generative phonology with weighted constraints, as in Harmonic Grammar (HG). For an overview of other approaches to gradience in phonology, see Pierrehumbert (2001).

3.2 Weighted constraints and gradient acceptability

Prince and Smolensky (1993/2004: 236) note that an OT-like theory of generative grammar could be constructed with HG's weighted constraints. In HG, the well-formedness of a linguistic representation is calculated by a Harmony function, which uses a linear equation of the same form as the one used in linear regression (see section 2.5) and other statistical procedures, and that is also made use of in connectionist networks (see Smolensky and Legendre 2006). Given a representation *R*, its scores on a set of constraints $\{C_1(R), C_2(R), C_3(R),..., C_n(R)\}$, and a set of coefficients, or weights, for those constraints $\{W_1, W_2, W_3,...,W_n\}$, Harmony is calculated by multiplying each score by the weight, and summing the weighted constraint scores, as in (6).

(6) $H(R) = W_1C_1(R) + W_2C_2(R) + W_3C_3(R) + ... W_nC_n(R)$

The optimal form in a candidate set can be chosen on the basis of Harmony scores: the optimum has maximal Harmony. Following Legendre *et al.* (2006), we convert OT's violation marks to negative integer scores, and following Keller (2006), we restrict constraint weights to non-negative real numbers (cf. the use of positive reals in Prince 2002 and Pater *et al.* 2007ab). An OT-style tableau that uses Harmony maximization as the criterion for optimality is provided in (7). Constraint weights appear in the top row, and the candidates' Harmony score appear in the rightmost column. Each candidate mapping violates one of the constraints. Because the first constraint has a higher weight, the candidate violating the second constraint has a higher Harmony score.

(7) *A weighted constraint tableau*

Weight	2	1	**H**
Input-1	Constraint 1	Constraint 2	
☞ Output 1-1		-1	-1
Output 1-2	-1		-2

One way that constraint weighting differs from ranking is in that it allows for cumulative constraint interactions, or 'gang effects'. This can be illustrated with the additional tableau in (8).

(8) *Cumulative interaction*

Weight	2	1	**H**
Input-2	Constraint 1	Constraint 2	
Output 2-1		-3	-3
☞ Output 2-2	-1		-2

This second tableau shows that three violations of Constraint 2 are worse than a single violation of Constraint 1, given the same constraint weighting as in (7). No OT ranking of these constraints could choose these two optima. If Constraint 1 were ranked above Constraint 2, as required to pick Output 1-1 in (7), then Output 2-1 would be optimal in (8).

Although Prince and Smolensky (1993/2004) and Legendre *et al*. (2006) claim that HG's cumulativity can produce implausible linguistic systems, the only published example of an unattested system produced by weighting but not ranking involves a gradient Alignment constraint (Legendre *et al*. 2006), which is controversial even in OT (see esp. McCarthy 2003a). The typological predictions of cumulative interaction are currently the subject of ongoing research (see Pater *et al*. 2007ab). Cumulativity has been claimed to be evidenced in well-formedness judgments in phonology (Ohala and Ohala 1986, Pierrehumbert 2001) as well as syntax (Keller 2006). See Pater (2007) for an HG account of the apparently greater prevalence of cumulative interactions in well-formedness judgments than in phonological alternations.

Another important difference between ranking and weighting is that the numerical score received by candidate mappings in HG can be used to model gradient acceptability. This is the original argument for HG in Legendre *et al*. (1990/2006), which is reiterated in Keller (2006). Harmony scores are used to model gradient and categorical phonotactics in Hayes and Wilson (2007), though as we discuss in section 5.2 the model of grammar they assume is somewhat different from the OT-like model we adopt.

To illustrate the HG account of gradient acceptability, we return to the case of gradient place co-occurrence restrictions in Arabic. All of the research reviewed in section 1, including our own descriptive analyses, agrees that there is at least a three-way distinction amongst coronals: pairs that agree in sonorancy and stricture (T-D) are the most restricted, with a near-zero O/E value, pairs that agree in neither sonorancy nor stricture (T-L) co-occur freely, while pairs that agree only in sonorancy (T-S) have an intermediate degree of attestedness. As discussed in the introduction, a theory with categorical OCP constraints faces a dilemma: if OCP-Coronal is relativized to continuancy, then T-S pairs are granted the same status as pairs such as T-L, and if not, then T-S pairs are ruled out. Although no experimental research seems to have focused specifically on these structures, Frisch and Zawaydeh's (2001) finding that acceptability of homorganic sequences in Arabic correlates negatively with similarity suggests that Arabic speakers would distinguish between these three types of forms.

Following Padgett (1995), we assume that there is a constraint that targets T-D but not T-S. In line with our descriptive analyses in section 2, we specify this constraint as applying between coronals with the same degree of stricture, so that it applies to pairs like [t-d] and [l-r], but not [s-t]. The constraint is defined in (9). We follow Suzuki's (1998) approach to OCP-type constraints in OT by defining constraints so that they apply between sequences of non-string-adjacent segments of a particular class. 'Sequence of consonants' in (9) refers to a consonant and the immediately following consonant, and ignores any intervening vowels.

(9) *T-T-STRIC
Assign a violation mark to a sequence of consonants in which both consonants are coronal and both consonants have the same stricture

We also assume a general constraint against sequences of coronals, and one against sequences of coronals that agree in sonorancy, as in (10) and (11).

(10) *T-T

Assign a violation mark to a sequence of consonants in which both consonants are coronal

(11) *T-T-SON

Assign a violation mark to a sequence of consonants in which both consonants are coronal and both consonants have the same specification for sonorancy

Since even sequences of coronals that agree in sonorancy are attested in Arabic, for all coronal sequences, the faithful candidate must be chosen over all unfaithful competitors. To illustrate, we will consider as an unfaithful candidate just the one in which the place specification of one of the consonants is altered, thus avoiding violations of the markedness constraints above. This candidate violates the faithfulness constraint in (12).

(12) IDENT-PLACE

Assign a violation mark to a consonant whose place specification is non-identical in Input and Output

In the following tableaux, we use the constraint values that were found in the learning simulation we report on in section 4. We continue to use labels for segment types: T and D stand for coronal oral stops, S stands for a coronal fricative, L a liquid, and P a labial stop. The weight of IDENT-PLACE is greater than the summed total of the markedness constraints, so that even violations of all three of them are preferred to a change in place specification. In terms of the choice of Input-Output mappings, the results for all three sequence types are equivalent: the faithful candidate always wins.

(13) Tableaux for Arabic sequence types

Weight	238.7	49.6	55.6	0	*H*
Input: TD	ID-PLACE	*T-T-SON	*T-T-STRIC	*TT	
☞ TD		-1	-1	-1	-115.2
PD	-1				-238.7

Weight	238.7	49.6	55.6	0	*H*
Input: TS	ID-PLACE	*T-T-SON	*T-T-STRIC	*TT	
☞ TS		-1		-1	-49.6
PS	-1				-238.7

Weight	238.7	49.6	55.6	0	*H*
Input: TL	ID-PLACE	*T-T-SON	*T-T-STRIC	*TT	
☞ TL				-1	0
PL	-1				-238.7

Although all three are equally optimal, these sequence types are distinguished in terms of their Harmony scores. TL is the most Harmonic, TD is the least, and TS receives an intermediate score. In these relative scores, we see the effect of HG cumulativity: TD is rated the lowest not because it violates any one constraint that is weighted the highest, but because it violates the most constraints. These scores correlate with the degree of attestedness of these sequence types in the Arabic lexicon (see section 4.4 on the degree of correlation), and we hypothesize that they would also correlate with Arabic speakers' judgments of these sequence types. In section 4.2, we provide an account of the acquisition of constraint weights such as these, and show how this grammar and learning model can be used to yield an account of hypothesized differences in Arabic and Muna speakers' knowledge of place restrictions. Before turning to the full analysis, we address an issue in the HG account of gradient acceptability we adopt here, and compare the resulting framework with some OT approaches.

3.3 Relativized Acceptability

In the analysis of Arabic phonotactics in the previous section, Harmony scores were used in two ways: to choose the optimal candidate in each tableau, and to assign relative acceptability to optimal candidates in different tableaux. In this respect, we follow Keller's (2006) previous proposals in syntax. Boersma (2004) points out a potentially fatal flaw for this approach. Because it does not relativize the acceptability score to a particular candidate set, in many cases suboptimal (ill-formed) structures will receive better scores than optimal (well-formed) ones. In the following schematic case, the optimal mapping /Input 1/ → [Output 1-2] receives a lower score than ungrammatical mapping /Input 2/ → [Output 2-1].

(14) *A problem for the Keller (2006) account of gradient acceptability*

Weight	3	2	1	**H**
/Input 1/	C1	C2	C3	
Output 1-1	-2			-6
☞ Output 1-2		-2		**-4**

Weight	3	2	1	**H**
/Input 2/	C1	C2	C3	
Output 2-1		-1		**-2**
☞ Output 2-2			-1	-1

The following tableaux provide a phonological example using an HG version of Lombardi's (1999) analysis of final devoicing (see Boersma 2004 for syntactic examples). In this analysis, voiced consonants, which violate *VOICE, are only permitted when devoicing would result in a violation of ID-ONS-VCE. This faithfulness constraint only applies in onset position. In coda position, only the general ID-VOICE constraint applies.

(15) *Ill-formed [pad] is more harmonic than well-formed [bam.bam]*

Weight	3	2	1	**H**
/pad/	ID-ONS-VCE	*VOICE	ID-VCE	
☞ [pat]			-1	-1
[pad]		-1		**-2**

Weight	3	2	1	**H**
/bambam/	ID-ONS-VCE	*VOICE	ID-VCE	
☞ [bam.bam]		-2		**-4**
[pam.pam]	-2		-2	-8
[bam.pam]	-1	-1	-1	-6

The ill-formed representation with a final voiced coda [pad] is given a Harmony score that is higher than that for the well-formed representation with a pair of onset voiced obstruents [bam.bam]. While this is not a problem for determination of optimality within the tableaux, it is an undesirable result if relative acceptability across tableaux is taken to be a function of relative Harmony scores, as in Keller's (2006) model. This problem is likely to arise in any case where a constraint can be violated only to satisfy another one; it is not specific to the interaction of particular types of constraint. To see that the problem is not specific to interactions with faithfulness constraints, we can take a case in which final stress is only permitted in monosyllabic words, due the domination of NONFINALITY by a constraint demanding that every word be stressed (see e.g. Prince and Smolensky 1993/2004: ch. 4). A well-formed monosyllable with final stress, and an ill-formed polysyllable with final stress, equally violate NONFINALITY. The source of the problem is very general: raw Harmony scores do not encode the relationship between a violation and the alternative structures that could avoid it.

To address this problem, we propose a revised acceptability metric for HG that makes the assignment of well-formedness scores relative to the candidate set. Under this definition, the acceptability of a representation is a function of the difference between its Harmony and the Harmony of its most harmonic competitor. For an ill-formed representation, the most harmonic competitor is the optimum, and for a well-formed representation, the most harmonic competitor is the next-best candidate.

(16) Acceptability(x) = H(x) - H(y)
Where y is the most harmonic candidate for the same input as x, and $y \neq x$.

Using this metric, well-formed structures get positive scores, and ill-formed ones get negative scores, as we can see by calculating the scores of the no longer problematic forms from the tableaux in (15).

(17) Acceptability([pad]) = H([pad]) - H([pat]) = (-2) - (-1) = -1
Acceptability([bambam]) = H([bambam]) - H([bampam]) = (-4) - (-6) = 2

For the Arabic example in the last section, the relative acceptability of the three consonant sequence types remains unchanged, since the competitor in each case receives the same score.

Relativized acceptability is not just a stipulated patch for Boersma's (2004) problem, as there is considerable precedent for the calculation of differences between the optimal candidate and other members of a candidate set. In OT, Prince and Smolensky's (1993/2004) Mark Cancellation yields the differences in candidates' violations, and Tesar and Smolensky's (2000) Constraint Demotion Algorithm operates on the basis of differences between optima and their competitors, as does Boersma's (1998) Gradual Learning Algorithm. As we will discuss in section 4.2, the HG learning algorithm that we adopt also uses differences between violations as the basis for updating constraint weights. Furthermore, the notion of a *margin of separation* between the correct outcome and the next-best alternative is central to many procedures in machine learning (see Boersma and Pater 2008 in the context of HG learning).

3.4 Comparison with OT approaches

There are a number of proposed extensions of standard OT that aim to capture various sorts of gradience. Most attention has been focused on variation between two or more acceptable outputs for a single input (e.g. Anttila 1997, Boersma and Hayes 2001, Coetzee 2006; see Coetzee and Pater 2008 for a review of approaches in OT and HG). Although there has been less research aimed at providing a formalization of gradient acceptability in OT, there are several extant proposals. Boersma and Hayes (2001) relate degree of acceptability to the probability of an output will be chosen when it is in variation with another output for the same input. This account does not extend to cases in which an output always surfaces faithfully, as in the Muna and Arabic cases we are considering, or to cases in which an output never surfaces faithfully, as in some cases of gradient ill-formedness. Hammond (2004) does apply a probabilistic version of OT to these sorts of cases, but does not provide an explanation for the absence of variation. Zuraw (2000) and Hayes and Londe (2006) provide an extension of Boersma and Hayes' model to cases of variation where particular morphemes may have fixed outcomes, but this only applies to morphophonology, and not to pure phonotactics.

Coetzee (2004, 2008) proposes a rank ordering account of gradient acceptability that is similar in some respects to the HG account discussed in section 3.2 (see also Everett and Berent 1998). The relative acceptability of two representations is determined by evaluating their relative performance on the constraint hierarchy. Because this comparison does not include any information about the other members of the candidate sets for each representation, it runs into the same problem that Boersma (2004) identifies for the Keller (2006) HG approach. If we take the final devoicing case as an example, the representation with a single voiced coda will be rated as better formed than one with a pair of onset voiced obstruents due to their relative perfomance on *VOICE. The fact that devoicing the onset would result in a violation of a higher ranked constraint has no bearing on the outcome.

(18) *A problem for OT rank-ordering*

Input	Output	ID-ONS-VCE	*VOICE	ID-VCE
/pad/	pad		*	
/bambam/	bambam		**	

There are OT accounts of gradient acceptability that calculate the acceptability of a representation with respect to candidate sets, but these account only for one side of the ill-formed/well-formed divide, and require considerable elaboration of the core theory. The account of gradient well-formedness in Pater (2005), which we applied to an earlier analysis of Muna

(Coetzee and Pater 2006), involves the postulation of lexically specific faithfulness constraints for every lexical item, and requires the calculation of the outcome of every possible indexation to determine degree of well-formedness for a nonce word. As Hayes and Wilson (2007) point out, it does not extend to gradient ill-formedness.[11] The account of gradient ill-formedness in Boersma (2004) involves the incorporation of numerical values of constraints, and the calculation of the percentage of times each candidate wins relative to the other candidates. This account does not extend to gradient well-formedness, since by definition attested (non-variable) forms win 100% of the time.

In sum, the numerically weighted constraints of HG allow for a straightforward account of gradient acceptability based on representations' Harmony scores. While OT's constraint ranking represents a departure from the categorical constraint satisfaction of other grammatical theories, there seems not yet to be a satisfactory account of gradient phonotactic acceptability available within that framework.

4. Muna and Arabic place restrictions in Harmonic Grammar

In section 2, we showed that Muna has a restriction against the co-occurrence of non-identical homorganic consonants that is gradient in two ways. First, the restriction is strongest amongst dorsals, somewhat weaker amongst labials, and weakest of all amongst coronals. Second, agreement of homorganic consonants in each of the features of sonorancy, voicing, stricture, and prenasalization results in lower rates of co-occurrence. In Muna, the contribution of each of these features is relatively equal, unlike Arabic, in which sonorancy agreement is the strongest predictor of attestedness of sequences of coronals. In section 3, we discussed evidence that speakers have knowledge of gradient phonotactics, and showed how the weighted constraints of HG can be used to represent such gradience.

In this section, we use HG weighted constraints to provide a hypothesized model of Muna and Arabic speakers' knowledge of the gradient phonotactics of their language. Section 4.1 provides the details of the constraint set. Section 4.2 introduces the learning algorithm we use to find a weighting for the constraints. Section 4.3 discusses the results obtained for Muna and section 4.4 continues the discussion of Arabic from section 3.2.

4.1 Constraint set

In section 3.2 we introduced a set of constraints against homorganic coronal sequences that we used to account for a subset of Arabic data. In this section, we provide the constraint set that we use in our full analysis of Arabic and Muna. First, we extend the set of constraints against homorganic sequences to other places of articulation, and relativize them for agreement for subsidiary features besides sonorancy and stricture. The features included are those that were shown to play a role in our descriptive analyses in section 2 (see Shiels-Djouadi 1975, Guy and Boberg 1997 and Côté 2004 on the role of voicing and other features in constraints on adjacent homorganic segments). In (19), we provide the schema for our constraints. As discussed above, these constraints apply only to contiguous sequences of consonants. The restrictions in Arabic

[11] Hayes and Wilson's (2007) diagnosis of the problem with the lexically specific constraint approach to gradient phonotactics is that the grammar is formed only on the basis of observed phonotactic distributions, and not with the help of a measure of expected distributions. However, in our current analysis, there is no calculation of an expected distribution, yet both gradient well-formedness and ill-formedness are accounted for. Whether an adequate theory of phonotactic learning requires explicit calculations of expected distributions remains an open question.

and Muna on 'non-adjacent' homorganic consonants are presumably due to more general constraints that do not have this restriction (see e.g. Suzuki 1998). We limit the constraints to non-identical consonants because identical consonants are not subject to any restriction in Muna. It might be preferable to derive the identity exemption from constraint interaction, but on the other hand, constraints against highly similar non-identical sequences may have a grounding in speech processing (see e.g. Hannson 2001, Frisch 2004). As discussed in section 2.1, we treat segments that differ only in prenasalization as identical.

(19) *Homorganic sequence constraint schema*

Assign a violation mark to a sequence of non-identical consonants that both have place of articulation *P* and agree in specification for *S*, where *P* ∈ {Pharyngeal, Dorsal, Coronal, Labial} and *S* ∈ {Sonorancy, Stricture, Voice, Emphatic, Prenasalization}

For the subsidiary features, we classified the segments as discussed in the descriptive analysis; see sections 2.2 to 2.4. The abbreviated names for the full expansion of this constraint schema appear in (20). P, T, K, H refer to labial, coronal, dorsal and pharyngeal place respectively. The abbreviations for the subsidiary features were chosen to transparently match the set of features in (19). We also include constraints that apply to sequences that share each place of articulation regardless of whether they agree for any other feature; their abbreviations lack a subsidiary feature specification (e.g. *P-P penalizes any sequence of non-identical labials).

(20) Constraints against homorganic sequences

*P-P	*P-P-SON	*P-P-STRIC	*P-P-VCE	*P-P-EMPH	*P-P-PRE
*T-T	*T-T-SON	*T-T-STRIC	*T-T-VCE	*T-T-EMPH	*T-T-PRE
*K-K	*K-K-SON	*K-K-STRIC	*K-K-VCE	*K-K-EMPH	*K-K-PRE
*H-H	*H-H-SON	*H-H-STRIC	*H-H-VCE	*H-H-EMPH	*H-H-PRE

In our analyses, we leave out constraints that target non-existent distinctions in each language. The Muna analysis does not include the constraints on pharyngeals or on emphatic agreement, and the Arabic analysis does not include the constraints on prenasal agreement, or on emphatic agreement in the non-coronals.

We also experimented with other, smaller constraint sets. For example, we tried to derive the difference between places of articulation from the cumulative interaction of constraints against each of the places of articulation (e.g. *DORSAL, which penalizes every dorsal) with more general OCP constraints that did not specify place of articulation. This approach failed dramatically, since words with a single dorsal or single labial are themselves rather well attested in the lexicon. Other constraint sets simply yielded less tight fits with the O/E values. It would also be possible to make the constraints more specific, for example relativizing them to agreement for a particular value of a feature (see Coetzee and Pater 2006). Ultimately, the correct degree of specificity of the constraints would likely be best assessed through a combination of typological work and experimental investigation of native speaker judgments.

4.2 Gradual learning and gradient phonotactics

As in OT, we assume that the constraint set is universal, and that phonological analysis, and phonological learning, requires finding an appropriate weighting (rather than ranking) of the constraints. In the analysis of categorical patterns we could follow similar logic as employed in OT. For example, to account for the absence of words with more than one prenasalized stop in Muna, it must be the case that the weight of a constraint against such words (*2-PRENAS) is greater than that of the faithfulness constraint penalizing denasalization between input and output (IDENT-PRE). The tableau in (21) illustrates the outcome produced by a weighting of *2-PRENAS above IDENT-PRE when an imaginary morpheme with two underlying prenasalized segments is submitted for evaluation: one of the prenasalized segments denasalizes.[12]

(21) *Denasalization in HG*

Weight	100	50	*H*
Input: ^{n}da^{m}pa	*2-PRENAS	IDENT-PRE	
☞ da^{m}pa		-1	-50
^{n}da^{m}pa	-1		-100

The constraint values in this tableau are partially arbitrary: any pair of values that respects the condition that *2-PRENAS is greater than IDENT-PRE will get the right outcome. For further discussion of the use of weighted constraints to account for categorical patterns in an OT-like version of HG, and of the use of logical conditions on weightings in such analysis, see Flemming (2001), Prince (2002), Legendre *et al*. (2006), and Pater *et al*. (2007ab).

To get constraint weights to produce results corresponding to gradient patterns, there is a wide range of potential analytic techniques available, given that the linear model of HG is broadly used in both statistical and connectionist modeling. In section 2.5, we provided a linear regression model that could be used to find weights for our full set of markedness constraints. There are two reasons why we do not apply that model here. First, our aim is to use a single grammar to account for both phonotactics and alternations, as in standard in generative phonology in general and OT in particular (in OT, see esp. McCarthy 2002, Hayes 2004, and Pater and Tessier 2006). Because the regression model does not incorporate interaction between faithfulness and markedness constraints, it would not meet that goal. Second, we do not regard a linear regression over observed and expected values to be a particularly realistic model of human language acquisition. Here, we adopt a simple learning algorithm that has been demonstrated to successfully find weights that generate both categorical patterns and ones with variation, and show that it yields weights that produce Harmony scores that correlate with gradient phonotactics.

The learning algorithm closely resembles the Gradual Learning Algorithm (GLA) for stochastic OT in Boersma (1998) and Boersma and Hayes (2001), and was first applied to phonology by Jäger (2008). Unlike the GLA for stochastic OT, which is demonstrably non-convergent (Pater 2008), this learning algorithm has proofs of convergence (Fischer 2005,

[12] The restriction on the co-occurrence of prenasalized segments is limited to the domain of the morpheme. Our Muna corpus does not contain a single morpheme with more than one such segment. However, affixes that contain prenasalized segments generally combine freely with roots that also contain prenasalized segments. The habitual affix /ma^{n}so/ combines freely with verbs like /li^{m}pu/ ‘ to forget’ to give a form like [no-ma^{n}so-li^{m}pu] ‘he is forgetful’ (van den Berg 1989:304). However, there is some evidence that a weaker form of this restriction applies even across morpheme boundaries. The 3rd plural possessive suffix has two allomorphs: /-do/ and /-^{n}do/. When attached to a root that contains a prenasalized segment, free variation is observed between the two forms as in [ka^{m}bele-^{n}do]~[ka^{m}bele-do] ‘their shadow’ (van den Berg 1989:36). However, with a root that does not contain a prenasalized consonant, only /-^{n}do/ is grammatical, as in [galu-^{n}do], *[galu-do] ‘their field’ (van den Berg 1989:85).

Boersma and Pater 2008). The learner is on-line and error-driven. It is supplied with input-output pairs, and uses the current state of its grammar to generate an output. Unlike the stochastic OT GLA, the learner uses an HG grammar of the form we have been assuming elsewhere in this paper. If the learner's own output diverges from that of the learning datum, learning is triggered. The weight update rule is one that is widely used in machine learning, where it is referred to as stochastic gradient ascent or the perceptron update rule. The update rule can be stated for HG grammars as in (22), where the learner's own output is termed the error, and the observed learning datum is the correct form.

(22) *Update rule for the GLA for HG*

Add n*(vE - vC) to each constraint weight

Where vE is the number of violations incurred by the error, and vC is the number of violations incurred by the correct form, and $0 < n \geq 1$

By increasing the weight of constraints that are violated more by the error, and decreasing the weight of constraints violated more by the correct form, the grammar is changed such that it is less likely to make this error when a similar piece of data is next encountered. The constant n determines the rate of change, and in some applications, can be adjusted over the course of learning (it is equivalent to the GLA's plasticity).

To illustrate how this update rule works, in (23) we show a tableau for an attested Muna sequence, with the constraint weights that we posit at the beginning of learning: the markedness constraints are all set at 100, and the faithfulness constraint at 50 (see Jesney and Tessier 2007 on the role of low faithfulness weighting in HG learning). With these weights, the learner produces the error [d-k]; the correct form is indicated with a check mark. The update rule will lower the value of the weights for all of the constraints violated by the correct form, and raise the value of all those violated by the error. The arrows beside the weight values indicate the direction of change for the constraints. The rate of change in the initial phase of the simulation is 1, so the markedness constraints would all be demoted to 99, and the faithfulness constraint would be raised to 51.

(23) A learning step

Weight	100 →	100 →	100 →	100 →	← 50	*H*
Input: g-k	*K-K-SON	*K-K-STRIC	*K-K-PRENAS	*K-K	ID-PLACE	
✓ g-k	-1	-1	-1	-1		-400.0
☞ d-k					-1	-50.0

This process will be repeated each time the learner is presented with the /g-k/~[g-k] datum, until the weight of the faithfulness constraint is sufficiently high that no further errors are made (see 4.3 for the final weighting values).

Our application of this approach to learning uses the implementation in Praat (Boersma and Weenink 2007). The evaluation mode was set to Linear OT, which limits constraint weights to positive reals, and automatically results in the update rule in (22) being applied (in versions of

Praat starting with 4.5.21). Settings were otherwise Praat's defaults, which includes noisy evaluation. As in stochastic OT, noisy evaluation can produce variation, but this is irrelevant to the outcome here, since the weights in the final state grammar are such that the output for each input does not vary across instances of evaluation.

The data for learning included all of the observed sequences of non-identical homorganic consonants from each language, distributed according to their frequency of occurrence in the lexicon. For example, the Muna lexicon we analyzed has no sequences of [ŋ-g], 2 of [k-g], 3 of [ŋ-k], and 15 of [ʁ-k]; our data file for Praat specified that they be presented to the learner in that relative distribution. In all cases, the input and output were simply identical, as is standardly assumed in OT phonotactic learning (see e.g. Hayes 2004, Prince and Tesar 2004).

The discussion of Arabic in section 3.2 illustrated part of the results of this learning simulation. In the next two sections we discuss in more detail the results, first for Muna, and then for Arabic.

4.3 Analysis of Muna

In (24), we give the constraint weights that our learner acquired for Muna. For the markedness constraints in this table, X-X should be replaced by K-K, P-P, or T-T to find the full constraint names.[13]

(24) Muna constraint weights

	K-K	P-P	T-T
*X-X	53.80	36.32	0
*X-X-SON	59.84	57.50	50.36
*X-X-STRIC	72.02	60.49	66.44
*X-X-VCE	96.83	80.54	52.87
*X-X-PRE	62.29	60.46	27.21
IDENT-PLACE		261.55	

From its starting position at 50, the faithfulness constraint IDENT-PLACE has been promoted to a value that is greater than the combined weights of any set of constraints that is violated by any existing sequence. Tableau (25) illustrates this for the [g-k] sequence whose initial state tableau appeared in (23). In all of our tableaux, the Harmony scores for the sequences are calculated from the actual weights; they occasionally differ slightly from the weighted sums in the tableaux here because of rounding.

(25) Tableau for attested [g-k]

Weight	261.6	59.8	72.0	96.8	62.3	53.8	*H*
Input: g-k	ID-PLACE	*K-K-SON	*K-K-STRIC	*K-K-VCE	*K-K-PRENAS	*K-K	
☞ g-k		-1	-1		-1	-1	-248.0
d-k	-1						-261.6

[13] The violation profiles that we used for each consonant pair in the learning simulation can be downloaded from www.umich.edu/~coetzee/Muna_Arabic/.

Some unattested sequences fail to be parsed faithfully by this grammar, as shown by the tableau in (26) for [ŋ-g].

(26) Tableau for unattested [ŋ-g]

Weight	261.6	59.8	72.0	96.8	62.3	53.8	*H*
Input: ŋ-g	ID-PLACE	*K-K-SON	*K-K-STRIC	*K-K-VCE	*K-K-PRENAS	*K-K	
ŋ-g			-1	-1	-1	-1	-284.9
☞ n-g	-1						-261.6

With this constraint weighting, the unattested [ŋ-g] sequence loses to the unfaithful candidate, and is thus categorized as ungrammatical or ill-formed. The difference between these outcomes is due to the greater weight of *K-K-VCE than *K-K-SON. There is only one attested sequence of voiced dorsals, [ᵑg-ŋ], which occurs only once in the corpus. This sole example of a violation of *K-K-VCE is thus presented very rarely to the learner (on average 1/1499 = 0.007% of the trials), and *K-K-VCE is demoted only to 96.8 before errors cease to be made on the [ᵑg-ŋ] sequence. On the other hand, the learner gets much more evidence for the demotion of *K-K-SON – there are 4 attested sequences that violate this constraint, namely the pairs of dorsal obstruents ([k-ᵑg], [k-g], [ᵑg-k], [ʁ-ᵑk]), and these sequences appear a total of 25 times in the corpus. These are thus presented to the learner with about 25 times the frequency of [ᵑg-ŋ] (mean 25/1499 = 1.7%), which results in *K-K-SON being demoted much further before learning ceases.

This illustrates the fact that frequency differences between structures are encoded in constraint weightings as an automatic byproduct of gradual learning (see also Zuraw 2000). The learner is doing no explicit calculation of the relative frequency of structures. We can see a further illustration of this effect, and its usefulness in accounting for gradient well-formedness, in the tableaux in (27).

(27) *Tableaux for Muna dorsal sequences*

Weight	261.6	59.8	72.0	96.8	62.3	53.8	*H*
Input: ŋ-k	ID-PLACE	*K-K-SON	*K-K-STRIC	*K-K-VCE	*K-K-PRENAS	*K-K	
☞ ŋ-k			-1		-1	-1	-188.1
n-k	-1						-261.6

Weight	261.6	59.8	72.0	96.8	62.3	53.8	*H*
Input: ʁ-k	ID-PLACE	*K-K-SON	*K-K-STRIC	*K-K-VCE	*K-K-PRENAS	*K-K	
☞ ʁ-k		-1			-1	-1	-175.9
s-k	-1						-261.6

In these tableaux, both sequences are parsed faithfully. Their Harmony scores differ, however, as a function of the difference between the weightings of *K-K-SON and *K-K-STRIC. As we mentioned just above, there are 4 *K-K-SON-violating sequences that appear 25 times in the corpus. There are also 4 sequences that violate *K-K-STRIC ([k-^{ŋ}g], [^{ŋ}g-ŋ], [k-ŋ], [k-g]), but there are only 8 occurrences of these sequences in the corpus. The learner is thus presented more than three times as many examples of *K-K-SON violations as *K-K-STRIC violations, and *K-K-STRIC winds up with a higher final weighting value.

As discussed in section 3.3, an acceptability score for candidate *x* is calculated by subtracting from the harmony score of *x* the harmony score of *x*'s most harmonic competitor *y* – i.e. Acceptability(*x*) = H (*x*) – H(*y*). For instance, the acceptability of the sequence [ŋ-k] from (27) above is calculated by subtracting the harmony score of its most harmonic competitor [n-k] from its own harmony score (Acceptability ([ŋ-k]) = H([ŋ-k]) - H([n-k]) = -188.1 - (-261.6) = 73.5). The acceptability scores of all the Muna sequences in the above tableau are provided in (28), along with their O/E values, calculated as discussed in section 2.

(28)

Sequence	*Acceptability*	*O/E*
ŋ-g	-23.3	0
k-g	13.6	0.06
ŋ-k	73.5	0.10
ʁ-k	85.7	0.40

Here we see that the acceptability scores correlate with the O/E scores; sequences that are better attested get higher acceptability scores. Because we lack data on Muna speakers' judgments of these forms, or other relevant experimental data, we follow Frisch *et al*. (2004) in taking it as a goal of our analysis of gradient phonotactics that the grammar categorize the forms in a way that correlates with O/E values. Some of the success of this analysis in that regard is illustrated in the correlation between acceptability and O/E values in (28); we report on the statistical analyses of these correlations below.

The data in (28) show how this analysis captures gradience due to differences in the featural makeup of segments within a single place of articulation. Cross-place differences are exemplified by the tableaux for the labial sequence type [m-p] and coronal [n-t] in (29), which can be compared with the dorsal [ŋ-k] in (27).

(29) *Cross-place differences illustrated*

Weight	261.6	57.5	60.5	80.5	60.5	36.3	*H*
Input: m-p	ID-PLACE	*P-P-SON	*P-P-STRIC	*P-P-VCE	*P-P-PRENAS	*P-P	
☞ m-p			-1		-1	-1	-157.3
n-p	-1						-261.6

Weight	261.6	50.4	66.4	52.9	27.2	0	*H*
Input: n-t	ID-PLACE	*T-T-SON	*T-T-STRIC	*T-T-VCE	*T-T-PRENAS	*T-T	
☞ n-t			-1		-1	-1	-93.6
m-t	-1						-261.6

The acceptability and O/E values for the cross-place comparison appear in (30). In this case, the differences between the sequences arise from the fact that the constraints against sequences of dorsals generally have the highest values, and those against sequences of coronals generally have the lowest, while the constraints against sequences of labials typically have an intermediate value. Again, the weighting differences arise due to the relative frequency of the sequences in the learning data.

(30)

Sequence	*Acceptability*	*O/E*
ŋ-k	73.5	0.10
m-p	104.3	0.39
n-t	168.0	0.70

While this analysis of gradient phonotactics reflects frequency differences, it is important to emphasize that it is not equivalent to an account based on raw segmental frequency. In our analysis, sequences with the same frequency will often receive different acceptability scores. This is illustrated clearly by the different scores assigned to unattested sequences, which all have the same frequency: 0. The tableaux for four further unattested dorsal sequences appear in (31).

(31) *Unattested dorsal sequences*

Weight	261.6	59.8	72.0	96.8	62.3	53.8	*H*
Input: g-ʁ	ID-PLACE	*K-K-SON	*K-K-STRIC	*K-K-VCE	*K-K-PRENAS	*K-K	
g- ʁ		-1		-1	-1	-1	-272.8
☞ d- ʁ	-1						-261.6

Weight	261.6	59.8	72.0	96.8	62.3	53.8	*H*
Input: ŋ-ᵑg	ID-PLACE	*K-K-SON	*K-K-STRIC	*K-K-VCE	*K-K-PRENAS	*K-K	
☞ ŋ-ᵑg			-1	-1		-1	-222.6
n-ᵑg	-1						-261.6

Weight	261.6	59.8	72.0	96.8	62.3	53.8	*H*
Input: g-ᵑk	ID-PLACE	*K-K-SON	*K-K-STRIC	*K-K-VCE	*K-K-PRENAS	*K-K	
☞ g-ᵑk		-1	-1			-1	-185.6
d-ᵑk	-1						-261.6

Weight	261.6	59.8	72.0	96.8	62.3	53.8	*H*
Input: ŋ-ᵑk	ID-PLACE	*K-K-SON	*K-K-STRIC	*K-K-VCE	*K-K-PRENAS	*K-K	
☞ ŋ-ᵑk			-1			-1	-125.8
n-ᵑk	-1						-261.6

The acceptability scores of the unattested dorsal sequences from the tableaux in (27) and (31) are provided in (32). Our analysis predicts that Muna speakers would distinguish these sequences, even though they are of equal frequency.

(32)

	Sequence	*Acceptability*	*O/E*
	ŋ-g	-23.3	0
	g-ʁ	-11.2	0
	ŋ-ᵑg	39	0
	g-ᵑk	76	0
	ŋ-ᵑk	135.8	0

This sort of prediction differentiates our approach, and other phonological accounts of gradient phonotactics, from a statistical account based purely on segmental frequency. In psycholinguistic experimentation, sequence frequency is typically measured in terms of the occurrence of the segments involved, without considering their featural makeup. Along with O/E, typical such measures are *n*-phone probability, transitional probability and neighborhood density (see Bailey and Hahn 2001 for an overview, and a proposal for a revised neighborhood density measure that makes use of Frisch *et al.*'s 2004 similarity metric).

There are two sources of predicted distinctions among the unattested sequences in (32), and these illustrate two broad classes of differences between phonological analyses of gradient phonotactics, and ones based on raw segmental frequency.[14] The first is *generalization* from the frequency of attested sequences. One example comes from the difference between the two sequence types that receive the lowest acceptability scores: [ŋ-g] and [g-ʁ]. Their violation profiles are identical except that [ŋ-g], but not [g-ʁ], violates *K-K-STRIC, while [g-ʁ], but not [ŋ-g], violates *K-K-SON. As discussed above, *K-K-SON is violated more frequently amongst the attested sequences, and thus winds up with a lower weight in the final state grammar. This distinction is then projected onto the unattested sequences in terms of their relative Harmony. Another example of this sort of effect is in the difference between [ŋ-ᵑg], which violates *K-K-VCE, and [g-ᵑk], which violates *K-K-SON.

The other source of distinctions between the unattested sequences derives from *universals* encoded in the constraint set. An example comes from the difference between [ŋ-ᵑk] and the other sequences in (32) containing prenasalized stops, [g-ᵑk] and [ŋ-ᵑg]. Sequences of the type [ŋ-ᵑk] violate a proper subset of the constraints violated by [g-ᵑk] and [ŋ-ᵑg]. Because of the additional constraint violations incurred by [g-ᵑk] and [ŋ-ᵑg], they are inherently less well-formed than [ŋ-ᵑk].

As discussed in section 3.1, there is experimental evidence showing that speakers do distinguish between consonantal sequences of zero frequency. Whether these distinctions are due to generalization from attested forms, to phonological universals, or to some other factor is an

[14] We note that these measures are usually applied in controlling stimuli, rather than in developing and testing a model of knowledge of phonotactics. Nonetheless, they provide a useful baseline for the development and testing of more elaborate phonological models. For further discussion of the relationship between these measures and phonological models, see Ohala and Ohala (1986), Coleman and Pierrehumbert (1997), Moreton (2002), Coetzee (2004, 2008, to appear), Myers and Tsay (2005), Albright (2007ab), Hammond (2007), Hayes and Wilson (2007), and Kirby and Yu (2007).

open question (see Berent *et al*. 2007 for recent discussion). We hope that with further development, the present model of gradient phonotactics will provide a framework for the further investigation of this issue.

It is also worth noting that only two of the five sequences in (32) lose to the unfaithful parse. It is clear from experimental work, as well as observations about loanword adaptation, that not all non-native sequences undergo repair, either in perception (see again Berent *et al*. 2007) or in production (e.g. Davidson 2006). Again, further development of this framework should allow us to model and predict such effects.

To determine the overall fit of the Harmony values to O/E, we examined the correlation of the O/E values to the acceptability scores, calculated as in (16). Since the most harmonic competitor of the winner is always the candidate that changes the place of one of the consonants, the acceptability score for all sequences are calculated by subtracting the weight of the IDENT-PLACE (261.6) from the Harmony score of the faithful sequence. The Harmony and acceptability scores are hence perfectly correlated for these data. Although our learning data did obviously not include any examples of consonant pairs with zero frequency, we do include these pairs in the regression in order to see whether our analysis can cope with these pairs. The r^2 value for this regression is 0.55 ($t(74) = 89.7$, $p < .001$), indicating a strong positive correlation between acceptability and O/E. Figure 1 shows the result of this regression analysis. We will return to this result in the comparison with Frisch *et al.*'s (2004) similarity metric in section 5.1.

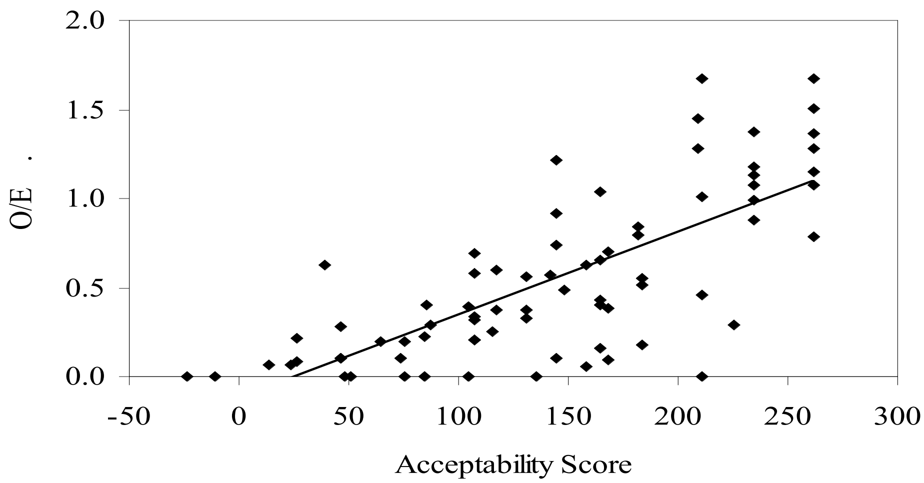

Figure 1: O/E vs. acceptability for Muna consonant pairs

4.4 Results for Arabic

In (33), we give the constraint weights that our learner acquired for Arabic; as for the Muna table, column and row labels can be combined to yield full constraint names.[15]

[15] The violation profiles that we used for each consonant pair in the learning simulation can be downloaded from www.umich.edu/~coetzee/Muna_Arabic/.

(33) *Arabic constraint weights*

	H-H	K-K	P-P	T-T
*X-X	56.69	67.14	100.00	0.00
*X-X-SON	73.47	73.39	100.00	65.74
*X-X-STRIC	74.66	76.32	100.00	69.17
*X-X-VCE	90.33	93.75	100.00	45.61
*X-X-EMPH				29.94
IDENT-PLACE	227.25			

Again, the weight of IDENT-PLACE is greater than the sum of the weights violated by any attested pair. As in Muna, we see place-related differences in the weights: because there are no labial pairs in the data presented to the learner, the constraints against labial sequences remained at their initial value.[16] Because there were relatively many coronal sequences, the constraints against them are weighted relatively low: the most general of these is at the minimum of zero imposed by Praat's LinearOT setting. Furthermore, we can note that within the coronals, the constraints that penalize sequences that agree in sonority and stricture have higher weights than those that penalize sequences that agree in voicing and emphatic, mirroring the relative frequency of these sequences in the data.

Because we have illustrated in detail how this sort of analysis works for Muna sequences with the tableaux in the previous section, and have illustrated the manner distinctions amongst Arabic coronals with tableaux in (13), we move directly to the regression analysis that measures the overall fit of acceptability scores to O/E for the Arabic pairs. Figure 2 depicts the result of performing the regression on all homorganic consonant pairs, excluding only the pairs of identical consonants. The r^2 value of this regression is 0.40 ($t(120) = 81.4$, $p < .001$). Figure 2 shows that there are a few consonant pairs in Arabic with very low acceptability scores. It is possible that these outliers are pulling r^2 down unduly. To check for this, we repeated the correlation analysis, excluding the six consonant pairs with negative acceptability scores. This resulted in an increase of 0.06 in r^2 to 0.46. As with Muna, there is a clear correlation between the acceptability scores and the O/E values.

[16] These weights are insufficiently high to guarantee that all labial sequences will be parsed unfaithfully: [bf], for example, which violates *P-P-SON and *P-P receives a score of -200, and is thus more harmonic than the unfaithful parse. As discussed in the previous section, it is not clear that all unattested sequences should lose to the unfaithful parse. However, this does illustrate the fact that gang effects between markedness constraints can push faithfulness above high default settings for other markedness constraints. This may eventually prove to force a different approach to biases and/or to phonotactic restrictiveness.

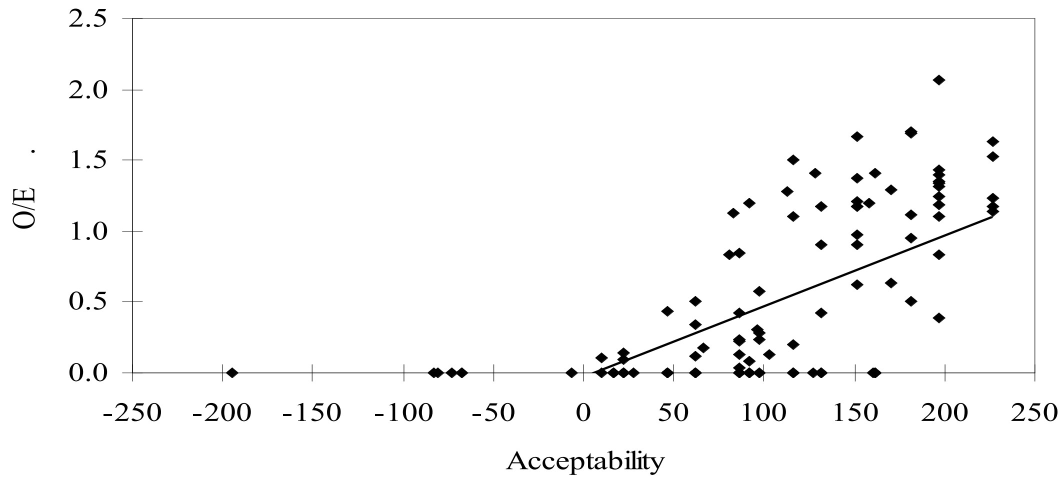

Figure 2: O/E vs. Acceptability for Arabic consonant pairs

5. Comparison with alternative accounts

5.1 Comparison with the similarity metric

In order to compare our account with the similarity avoidance account of Frisch *et al*. (2004), we calculated similarity values for all homorganic consonant pairs in Muna according to Frisch *et al*.'s natural classes similarity metric. The feature set that we used is included in Appendix A. As far as possible, we followed the same principles as Frisch *et al*. (2004) did when they calculated the similarity scores of Arabic consonant pairs. The feature set that we used also agrees with the assumptions that we made about Muna consonants in section 2 above. We used Adam Albright's *Segmental Similarity Calculator* to calculate the similarity of all the homorganic consonant pairs. The results of this calculation are included in Appendix B.[17]

We then performed a regression analysis to determine the correlation between the similarity values and the O/E values for Muna. The result of this comparison is shown in Figure 3, and should be compared with Figure 1 above where the correlation between Harmony and O/E values was plotted. We performed the same regression analysis for Arabic, using the similarity values reported by Frisch *et al*. (2004). The results for Arabic are represented in Figure 4, and should be compared the Arabic Harmony results in Figure 2 above.

[17] The *Segmental Similarity Calculator* can be downloaded from http://web.mit.edu/albright/www/. The feature files that we used as input to the *Segmental Similarity Calculator* as well the output generated by the software are available from the authors or can be downloaded from www.umich.edu/~coetzee/Muna_Arabic/.

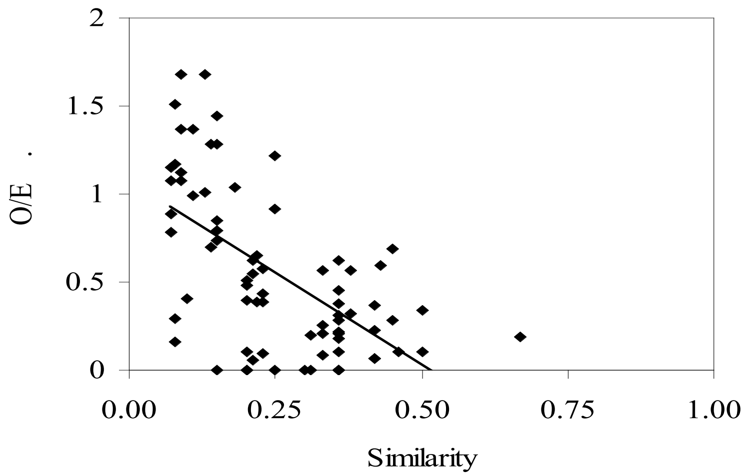

Figure 3: O/E vs. Similarity for Muna consonant pairs

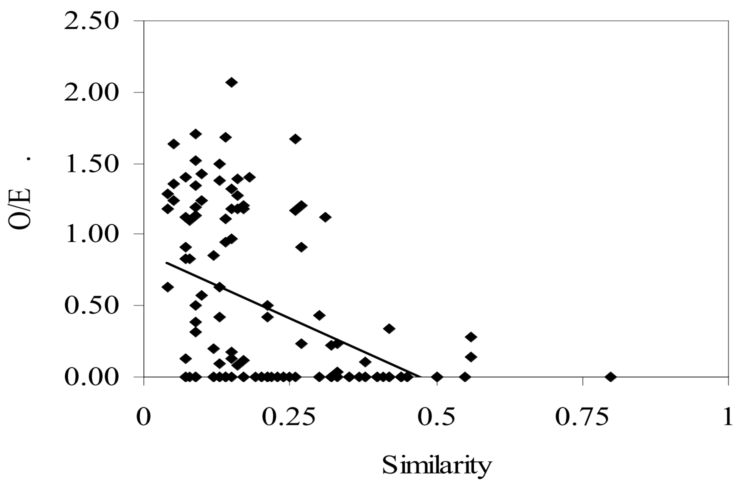

Figure 4: O/E vs. Similarity for Arabic consonant pairs

Inspection of Figures 3 and 4 shows clear evidence for a negative correlation between similarity and O/E, as is expected based on the results of Frisch *et al*. (2004). To compare the similarity-based account with the Harmony-based one, Table 14 contains the r^2 values for several different regressions performed on the O/E data of Muna and Arabic. The first row gives the r^2 values for regression analyses performed on all homorganic non-identical pairs. These are the regressions depicted in Figures 1 to 4 above. The second line includes all homorganic pairs in Arabic, identical as well as non-identical. To generate the Harmony scores for this comparison, we ran a learning simulation identical to the one reported above except that we included the one observed identical pair ([d-d]) in the learning data, and added a constraint against identical pairs. In calculating the Harmony scores for the identical pairs, we assumed that they also violated the constraints against the homorganic sequences in (20). The final line shows the results of performing a regression only on the non-identical coronal pairs of Arabic. It does not matter how the regression is performed or how the data are partitioned, the fit between the Harmony scores and the O/E values is better than that between the Similarity score and the O/E values.[18]

	Muna		*Arabic*	
	Harmony	*Similarity*	*Harmony*	*Similarity*
All non-identical pairs	0.55	0.36	0.40	0.20
Identical pairs included			0.29	0.21
All coronal pairs			0.45	0.25

Table 14: Comparing r^2 values of the Similarity account and the Harmonic Grammar account of Muna and Arabic consonant co-occurrence restrictions

It is perhaps not surprising that the Harmony-based account outperforms the similarity metric, since our learning simulation essentially fits the constraint values to the observed data. This is in line with one of the central points of our paper: that the cross-linguistic differences found in place co-occurrence patterns go beyond what is predicted by the similarity metric. Given such differences, an account of them must provide learners with a means of constructing language-specific grammars. An account with universal constraints and language-specific weightings allows us to capture similarities and differences in the Muna and Arabic patterns; our approach to learning provides a simple account of how these grammars may be acquired.

The similarity metric does provide one mechanism for deriving differences in place co-occurrence patterns: the structure of the segmental inventory. For Arabic, Frisch *et al*. (2004) claim that the weakness of the restriction amongst coronals is due to the large size of the Arabic coronal inventory. Arabic has 14 coronals, 6 pharyngeals, 3 dorsals and 3 labials. This affects the calculation of similarity by increasing the set of natural classes that each segment in a pair participates in, thus increasing the denominator in the equation in (2). The differences across places of articulation in Muna provide evidence that the relative strength of place co-occurrence patterns is not due to inventory structure alone. The average similarity and O/E values for the pairs that can be compared across all three places of articulation (voiced and voiceless stops, nasals, and prenasalized stops), omitting those that may be subject to the identity effect or the ban on multiple prenasalized stops, are given in Table 15.

	Coronals	*Labials*	*Dorsals*
O/E	0.78	0.22	0.09
Similarity	0.26	0.30	0.35

Table 15 : Muna O/E and average similarity values

[18] Frisch *et al*. (2004) report a much higher r^2 value for the relationship between similarity and O/E for the coronal pairs of Arabic (0.75). Their regression differed in several important ways from ours, which makes a direct comparison difficult. First, they group consonant pairs according to their similarity values. As a result, they have fewer data points to fit, which could lead to a higher correlation. Secondly, they do not perform a linear regression but rather use a stepwise monotonically decreasing regression function. Later in their paper (p. 209), they show that this function is practically equivalent to a logarithmic regression. If we perform a logarithmic regression without grouping the coronals into similarity classes, we get a somewhat higher r^2 value of 0.27. However, this value is still considerably lower then the r^2 that we obtain for the ordinary linear regression using Harmony scores. We therefore conclude that the Harmony scores match up better with the O/E values than do the similarity values.

The differences amongst the similarity values are relatively small; these would all be placed together in Frisch *et al*.'s (2004) groupings (see their Table 4). The O/E values, however, fall dramatically from the coronals to the labials. This problem appears to be due to the fact that the coronal inventory size in Muna is only marginally larger than that of the labials: Muna has 10 coronals, 8 labials and 6 dorsals. McCarthy (2003b) finds a similar problem in Rotuman. Like Arabic, co-occurrence between coronals in Rotuman is much freer across sonorancy classes than within them, but unlike Arabic, the coronal inventory is not particularly large. The Rotuman and Muna facts suggest that differences in co-occurrence restrictions across places of articulation are not due to inventory size alone.

5.2 Comparison with other Harmony-based accounts of gradient well-formedness

In this section, we comment briefly on the relationship of our analysis of gradient phonotactics to other accounts of gradient well-formedness that use a Harmony-based measure.

The original account of gradient acceptability in HG is Legendre *et al*.'s (1990) treatment of syntactic gradience. Our approach differs not only in that it focuses on gradient phonotactics, but also that it is cast within a grammatical theory that makes use of OT's notion of optimality in a candidate set. In that respect, a closer precedent is Keller's (2006) work on syntactic gradience, which also uses Harmony scores both to choose the optimal form in a single tableau, and to measure degree of well-formedness across tableaux. Section 3.3 showed how a problem with this approach raised by Boersma (2004) can be addressed by making gradient acceptability relative to the candidate set. Another difference in our approach is in how the constraint weights are set. Keller (2006) uses linear regression to achieve a correlation between Harmony scores and judgment scores; the judgment scores are taken as given. In contrast, we used a learning algorithm, and the observed data, to generate predicted acceptability scores. An obvious direction for future research, in both phonology and syntax, is to combine these approaches: to test whether predictions from grammars learned on the basis of observed data do match judgments.

This approach is taken in some of Hayes and Wilson's (2007) recent work (see also Albright 2007ab), whose model resembles ours in its use of the Harmony function to predict acceptability. Their model differs from ours in a number of other ways, however. First, it provides a mechanism for generating constraints, while we have assumed a universal constraint set. As mentioned in section 4.3, the question of whether phonotactic knowledge reflects universals encoded in a constraint set, or can be derived from the observed forms in a language, is currently an open one. There are also many ways to account for cross-linguistic universals and language-specific restrictions in phonotactics besides ranking or weighting a universal phonological constraint set, and we look forward to future research that develops and compares various alternatives (see Boersma 2007 and Hayes and Wilson 2007 for recent discussion).

Another way in which the Hayes and Wilson (2007) proposal differs from ours is in that their model does not perform a mapping from input to output: it directly assigns scores to surface forms. This has two consequences. One is that phonotactics and alternations are given separate accounts, unlike standard OT. Since our model retains this aspect of OT, it inherits OT's resolution of the duplication problem (Kenstowicz and Kisseberth 1977; McCarthy 2002). The other consequence is that because there is no candidate set in the Hayes and Wilson model, it cannot make use of the central OT concept of resolving constraint conflict through prioritization: its constraints are surface-true. Hayes and Wilson point out that a Maximum Entropy model like theirs can be used to calculate the probability of outputs conditional on a given input, thus yielding a theory closer to standard OT. Such a theory would be very close to the one we adopt; see the discussion of "MaxEnt-HG" and "Noisy-HG" in Coetzee and Pater (2008). A better understanding of the differences between these and other models of phonotactics will require further modeling of experimental results and typological data, as well as closer consideration of the subset problem for phonotactic learning (Smolensky 1996, Hayes 2004, Prince and Tesar 2004); see Jesney and Tessier (2007) for initial results on the subset problem using an approach similar to ours.

6. Conclusions

We have shown that Muna displays consonantal place co-occurrence restrictions that resemble those in Arabic in that their strength depends both on the place of articulation of the consonants, and on whether the consonants agree in their specification for other features. The languages differ, however, in the relative contribution of the subsidiary features. While sonorancy agreement is the prime predictor of attestedness of homorganic pairs in Arabic, in Muna there is a more balanced contribution across these features, which also include voicing and stricture, and to a lesser extent, prenasalization.

Based on experimental work that shows that speakers display gradient knowledge of consonantal phonotactics, we hypothesized that Muna and Arabic speakers have knowledge of these gradient place restrictions. We then showed how this knowledge can be modeled in terms of the weighted constraints of HG, and provided learning simulations for Arabic and Muna that found sets of weights that yield acceptability scores that correlate with O/E values. This correlation is stronger than the correlation between O/E and scores derived from Frisch *et al.*'s (2004) similarity metric.

Much remains to be done in terms of further developing this model of gradient phonotactic acceptability, and testing its predictions against experimental data. However, there are some general theoretical conclusions that we can draw at this point. Models of variation in OT such as those of Anttila (1997), Boersma and Hayes (2001) and Coetzee (2006) show that the scope of the theory can be extended from purely categorical patterns to cover certain kinds of gradience. This development can only be seen as positive, since with a given constraint set, these theories generate exactly the same categorical patterns as standard OT; the extension of the theory to cover variation does not yield an overly powerful theory of categorical phonology. Our work, and that of other proposals in HG, aims to further extend the theory to types of gradience that do not appear to be properly handled by ranked constraints, by replacing ranking with weighting (see also section 4.2 for discussion of advantages for the modeling of learning variable patterns). Because we maintain all other aspects of standard OT's architecture, this theory continues to generate all of the categorical patterns that OT does. As discussed in section 3.4, whether it overgenerates can only be determined through further research. However, we see this as a risk worth taking, since we believe that a theory of phonology that goes beyond a two-way distinction in acceptability moves toward being a more realistic model of phonological knowledge.

References

Albright, Adam. 2007a. *Gradient Phonological Acceptability As a Grammatical Effect*, ms., MIT. (http://web.mit.edu/albright/www/papers/Albright-GrammaticalGradience.pdf)

Albright, Adam. 2007b. *Natural Classes Are Not Enough: Biased Generalization in Novel Onset Clusters*, ms., MIT. (http://web.mit.edu/albright/www/papers/AlbrightBiasedGeneralization.pdf)

Anttila, Arto. 1997. 'Deriving Variation from Grammar', in F. Hinskens, R van Hout & W. L. Wetzels (eds.), *Variation, Change and Phonological Theory*, John Benjamins, Amsterdam, pp. 35-68).

Bailey, Todd and Ulrike Hahn. 2001. 'Determinants of Wordlikeness: Phonotactics or Lexical Neighborhoods?', *Journal of Memory and Language* **44**, 568–591.

Berent, Iris, and Joseph Shimron. 1997. 'The Representation of Hebrew Words: Evidence from the Obligatory Contour Principle', *Cognition* **64**, 39-72.

Berent, Iris, Daniel L. Everett, and Joseph Shimron. 2001. 'Do Phonological Representations Specify Variables? Evidence from the Obligatory Contour Principle', *Cognitive Psychology* **42**, 1-60.

Berent, Iris, Donca Steriade, Tracy Lennertz, and Vered Vaknin. 2007. 'What We Know About What We Have Never Heard: Evidence From Perceptual Illusions', *Cognition* **104**, 591–630.

Berkley, Deboran. 2000. *Gradient Obligatory Contour Principle Effect*, unpublished Ph.D. dissertation, Northwestern University, Evanston.

Boersma, Paul. 1998. *Functional Phonology: Formalizing the Interactions between Articulatory and Perceptual Drives*, Ph.D dissertation, University of Amsterdam.

Boersma, Paul. 2004. *A Stochastic OT Account of Paralinguistic Tasks Such as Grammaticality and Prototypicality Judgments*, ms., University of Amsterdam. (ROA-648).

Boersma, Paul 2007. *The Evolution of Phonotactic Distributions in the Lexicon*, paper presented at the Stanford Workshop on Variation, Gradience and Frequency in Phonology in July 2007, slides available at http://www.stanford.edu/dept/linguistics/linginst/nsf-workshop/BoersmaSlides.pdf.

Boersma, Paul, and Bruce Hayes. 2001. 'Empirical Tests of the Gradual Learning Algorithm', *Linguistic Inquiry* **32**, 45-86.

Boersma, Paul, and Joe Pater. 2008. *Convergence Properties of a Gradual Learner in Harmonic Grammar*, ms., University of Amsterdam and University of Massachusetts Amherst.

Boersma, Paul, and David Weenink. 2007. *Praat: Doing Phonetics by Computer*, computer program developed at the Institute of Phonetic Sciences, University of Amsterdam (http://www.praat.org/).

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

Coetzee, Andries W. 2004. *What It Means to Be a Loser: Non-Optimal Candidates in Optimality Theory*, unpublished Ph.D. dissertation, University of Massachusetts Amherst.

Coetzee, Andries W. 2006. ‘Variation as Accessing Non-Optimal Candidates’, *Phonology* **23**, 337-385.

Coetzee, Andries W. 2008. ‘Grammar is Both Categorical and Gradient’, in Steve Parker (ed.), *Phonological Argumentation*, Equinox, London. (In press.)

Coetzee, Andries W. to appear. ‘Grammaticality and ungrammaticality in phonology’, *Language*.

Coetzee, Andries W. and Joe Pater. 2006. ‘Lexically ranked OCP-Place constraints in Muna’, ms, University of Michigan and University of Massachusetts, Amherst (ROA-842).

Coetzee, Andries W. and Joe Pater. 2008. ‘The place of variation in phonological theory’, ms, University of Michigan and University of Massachusetts, Amherst (ROA-946).

Cohn, Abigail. 1992. ‘The Consequences of Dissimilation in Sundanese’, *Phonology* **9**, 199-220.

Coleman, Jennifer. S., and Janet Pierrehumbert. 1997. ‘Stochastic Phonological Grammars and Acceptability’, in *Computational Phonology. Third Meeting of the ACL Special Interest Group in Computational Phonology*, Association for Computational Linguistics, Somerset, NJ, pp. 49–56.

Côté, Marie-Hélène. 2004. ‘Syntagmatic Distinctness in Consonant Deletion’, *Phonology* **21**, 1-41.

Cowan, J. 1979. *Hans Wehr: A Dictionary of Modern Written Arabic*, Otto Harrasowitz, Wiesbaden.

Davidson, Lisa. 2006. ‘Phonology, Phonetics, or Frequency: Influences on the Production of Non-Native Sequences’, *Journal of Phonetics* **34**, 104-137.

Dresher. B. Elan. 1989. *Comments on Guttural Phonology*, paper presented at MIT Conference on Feature Geometry.

Evans, N. 1995. ‘Current Issues in the Phonology of Australian Languages’, in J. Goldsmith (ed.), *The Handbook of Phonological Theory*, Blackwell, Cambridge, MA.

Everett, Daniel, and Iris Berent. 1998. *The Comparative Optimality of Hebrew Roots: An Experimental Approach to Violable Identity Constraints*, ms., University of Pittsburgh and Florida Atlantic University. (ROA-235)

Faraway, Julian J. 2005. *Linear Models with R*, Chapman & Hall, New York.

Fischer, Markus. 2005. *A Robbins-Monro type learning algorithm for an entropy maximizing version of stochastic Optimality Theory*. MA thesis, Humboldt University, Berlin. (ROA-767).

Flemming, Edward. 2001. ‘Scalar and Categorical Phenomena in a Unified Model of Phonetics and Phonology’, *Phonology* **18**, 7-44.

Frisch, Stefan A. 2004. Language processing and segmental OCP effects. In Bruce Hayes, Robert Kirchner and Donca Steriade (eds.) *Phonetically Based Phonology*. Cambridge University Press. 346-371.

Frisch, Stefan A., and Bushra Adnan Zawaydeh. 2001. 'The Psychological Reality of OCP-Place in Arabic', *Language* **77**, 91-106.

Frisch, Stefan A, Janet Pierrehumbert, and Michael Broe. 2004. 'Similarity Avoidance and the OCP', *Natural Language and Linguistic Theory* **22**, 179-228.

Greenberg, J. H. 1950. 'The Patterning of Root Morphemes in Semitic', *Word* **6**, 162-181.

Greenberg, J.H. and J. Jenkins. 1964. 'Studies in the Psychological Correlates of Sound System of American English', *Word* **20**, 157-177.

Guy, Gregory R., and Charles Boberg. 1997. 'Inherent Variability and the Obligatory Contour Principle', *Language Variation and Change* **9**, 149–164.

Hammond, Michael. 2004. 'Gradience, Phonotactics, and the Lexicon in English Phonology', *International Journal of English Studies* **4**, 1-24.

Hammond, Michael. 2007. *Typology, Judgments and Weights*, paper presented at the Stanford Workshop on Variation, Gradience and Frequency in Phonology in July 2007, slides available at http://www.stanford.edu/dept/linguistics/linginst/nsf-workshop/HammondSlides .pdf

Hannson, Gunnar. 2001. 'The phonologization of production constraints: Evidence from consonant harmony', in M. Andronis, C. Ball, H. Elston and S. Neuvel (eds.) *CLS 37: The Main Session. Papers from the 37th Meeting of the Chicago Linguistic Society, Vol. 1*, Chicago Linguistic Society, Chicago. pp. 187-200.

Hayes, Bruce. 2004. 'Phonological acquisition in Optimality Theory: the early stages', in René Kager, Joe Pater and Wim Zonneveld (eds.) *Constraints in Phonological Acquisition. Cambridge*, Cambridge University Press. pp. 158-203.

Hayes, Bruce, and Zsuzsa Cziráky Londe. 2006. 'Stochastic Phonological Knowledge: The Case of Hungarian Vowel Harmony', *Phonology* **23**, 59-106.

Hayes, Bruce, and Colin Wilson. 2007. 'A Maximum Entropy Model of Phonotactics and Phonotactic Learning', *Linguistic Inquiry*. (In press.)

Itô, Junko and Armin Mester. 1986. 'The phonology of voicing in Japanese: Theoretical consequences for morphological accessibility' *Linguistic Inquiry* **17**, 49-73.

Jäger, Gerhard and Anette Rosenbach. 2006. 'The winner takes it all - almost. Cumulativity in grammatical variation', *Linguistics* **44**, 937-971.

Jäger, Gerhard. 2008. 'Maximum Entropy Models and Stochastic Optimality Theory', in Jane Grimshaw, Joan Maling, Chris Manning, Jane Simpson, and Annie Zaenen (eds.), *Architectures, Rules, and Preferences: A Festschrift for Joan Bresnan*, CSLI Publications, Stanford. (In press.)

Jesney, Karen and Anne-Michelle Tessier. 2007. 'Re-evaluating learning biases in Harmonic Grammar', in M. Becker (ed.), *University of Massachusetts Occasional Papers 36: Papers in Computational Phonology*. GLSA, Amherst.

Kawahara, Shigeto. 2007. 'Half Rhymes in Japanese Rap Lyrics and Knowledge of Similarity', *Journal of East Asian Linguistics* **16**, 113-144.

Kawahara, Shigeto, Hajime Ono, and Kiyoshi Sudo. 2006. ‘Consonant Co-occurrence Restrictions in Yamato Japanese’, *Japanese/Korean Linguistics* **14**, 27-38.

Keller, Frank. 2006. ‘Linear Optimality Theory as a Model of Gradience in Grammar’, in Gisbert Fanselow, Caroline Féry, Ralph Vogel, and Matthias Schlesewsky, (eds.), *Gradience in Grammar: Generative Perspectives*. Oxford: Oxford University Press. (In press.)

Kenstowicz, Michael and Charles Kisseberth. 1977. *Topics in Phonological Theory*. Academic Press, NewYork.

Kirby, James P. and Alan C.L. Yu. 2007. ‘Lexical and phonotactic effects on wordlikeness judgments in Cantonese’, in Jürgen Trouvain andWilliam J. Barry, (eds.) *ICPhS XVI*: *Proceedings of the 16th International Congress of Phonetics Science*. pp. 1389-1392.

Lombardi, Linda. 1999. Positional Faithfulness and Voicing Assimilation in Optimality Theory. *Natural Language and Linguistic Theory* **17**. 267-302.

Leben, William. 1973. *Suprasegmental Phonology*, unpublished Ph.D. dissertation, MIT, Cambridge, MA.

Legendre, Géraldine, Yoshiro Miyata, and Paul Smolensky. 1990. ‘Harmonic Grammar – a Formal Multi-Level Connectionist Theory of Linguistic Wellformedness: An Application’, in *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, Lawrence Erlbaum, Cambridge, MA, pp. 884–891.

Legendre, Géraldine, Antonella Sorace, and Paul Smolensky. 2006. ‘The Optimality Theory–Harmonic Grammar Connection’, in Paul Smolensky and Géraldine Legendre (eds.), *The Harmonic Mind: From Neural Computation to Optimality-Theoretic Grammar*, MIT Press, Cambridge, MA, pp. 903–966.

MacEachern, Margaret R. 1999. *Laryngeal Cooccurrence Restrictions*, Garland, New York.

Mann, H. B., and D.R. Whitney. 1947. ‘On a Test of Whether One of Two Random Variables is Stochastically Larger Than the Other’, *Annals of Mathematical Statistics* **18**, 50-60.

McCarthy, John J. 1988. ‘Feature Geometry and Dependency: A Review’, *Phonetica* **45**, 84–108.

McCarthy, John J. 1994. ‘The Phonetics and Phonology of Semitic Pharyngeals’, in Patricia Keating (ed.), *Phonological Structure and Phonetic Form: Papers in Laboratory Phonology III*, Cambridge University Press, Cambridge, pp. 191-233.

McCarthy, John J. 2002. *A Thematic Guide to Optimality Theory*. Cambridge University Press, Cambridge.

McCarthy, John J. 2003a ‘OT Constraints Are Categorical’, *Phonology* **20**, 75-138.

McCarthy, John J. 2003b. *Rotuman Consonant Cooccurrence Restrictions*, Course Handout, Ling 730, University of Massachusetts Amherst.

McConvell, P. 1988. ‘Nasal Cluster Dissimilation and Constraints on Phonological Variables in Gurindji and Related Languages’, in N. Evans and S. Johnson (eds.), *Aboriginal Linguistics 1*, Department of Linguistics, University of New England, Armidale.

Mester, Armin. 1986. *Studies in Tier Structure*, unpublished Ph.D. dissertation, University of Massachusetts Amherst.

Miles, Jeremy. 2005. 'Tolerance and Variance Inflation Factor", in Brian S. Everitt and David C. Howell (eds.), *Encyclopedia of Statistics in Behavioral Science*, John Wiley & Sons, Hoboken, N.J., pp. 2055-2056.

Moreton, Elliott. 2002. 'Structural Constraints in the Perception of English Stop-Sonorant Clusters', *Cognition* **84**, 55-71.

Myers, James, and Jane Tsay. 2005. 'The Processing of Phonological Acceptability Judgments', in *Proceedings of Symposium on 90-92 NSC Projects*, Taipei, Taiwan, pp. 26-45.

Ohala, John J. and Manjari Ohala. 1986. 'Testing Hypotheses Regarding the Psychological Manifestation of Morpheme Structure Constraints' in John J. Ohala, and Jerri Jaeger (eds.), *Experimental Phonology*, Academic Press, Orlando, pp. 239-252.

Padgett, Jaye. 1995. *Stricture in Feature Geometry*, Center for the Study of Language and Information, Stanford University.

Pater, Joe. 2005. 'Learning a Stratified Grammar', in A. Brugos, M. R. Clark-Cotton, and S. Ha (eds.), *Proceedings of the Boston University Conference on Language Development*, Cascadilla Press, Somerville, MA, pp. 482-492.

Pater, Joe. 2007. 'Cumulative ill-formedness in typological and experimental data', Paper presented at the Conference on Experimental Approaches to Optimality Theory, University of Michigan, handout available at http://people.umass.edu/pater/pater-michigan.pdf.

Pater, Joe. 2008. 'Gradual Learning and Convergence', *Linguistic Inquiry* 39/2.

Pater, Joe, Rajesh Bhatt, Chris Potts. 2007a. *Linguistic Optimization*, ms., University of Massachusetts Amherst.

Pater, Joe, Chris Potts, Rajesh Bhatt. 2007b. *Harmonic Grammar with Linear Programming*, ms., University of Massachusetts Amherst.

Pater, Joe, and Anne-Michelle Tessier. 2006. 'L1 Phonotactic Knowledge and the L2 Acquisition of Alternations', In R. Slabakova, S. Montrul, and P. Prévost (eds.), *Inquiries in Linguistic Development: Studies in Honor of Lydia White*. John Benjamins Publishing Company, Amsterdam/Philadelphia. 115-131.

Peperkamp, Sharon. 2007. 'Do We Have Innate Knowledge about Phonological Markedness?' Comments on Berent, Steriade, Lennertz, and Vaknin', *Cognition* **104**, 631-637.

Pertz, D. L., and T.G. Bever. 1975. 'Sensitivity to Phonological Universals in Children and Adolescents', *Language* **51**, 149-162.

Pierrehumbert, Janet. 1993. 'Dissimilarity in the Arabic Verbal Roots', in A. Schafer (ed.), *NELS 23: Proceedings of the North East Linguistic Society*, GLSA, Amherst, pp. 367-381.

Pierrehumbert, Janet. 2001. 'Stochastic Phonology', *GLOT* **5**, 1-13.

Prentice, D.J. 1971. *The Murut Languages of Sabah*. Australian National University, Canberra.

Prince, Alan S. 2002. 'Anything Goes', in T. Honma, M. Okazaki, T. Tabata, and S. Tanaka (eds.), *New Century of Phonology and Phonological Theory*, Kaitakusha, Tokyo, pp. 66–90.

Prince, Alan S., and Paul Smolensky. 1993. *Optimality Theory: Constraint Interaction in Generative Grammar*, ms., Rutgers University Center for Cognitive Science, New Brunswick, NJ. [Published as Prince and Smolensky 2004.]

Prince, Alan S., and Paul Smolensky. 2004. *Optimality Theory: Constraint Interaction in Generative Grammar*, The MIT Press, Cambdridge, MA. [Published version of Prince and Smolensky 1993.]

Prince, Alan and Bruce Tesar. 2004. Learning phonotactic distributions, in René Kager, Joe Pater and Wim Zonneveld (eds.) *Constraint in Phonological Acquisition*. Cambridge, Cambridge University Press. pp. 245-291.

Scholes, R. 1966. *Phonotactic Grammaticality*, Mouton, The Hague.

Selkirk, Elisabeth. 1991. *Vowel Height Features: Evidence for Privativity and Dependency*, paper presented at UQAM, Montréal.

Shiels-Djouadi, M. 1975. 'Reappraisal of the Voicing Constraint in Consonant Cluster Simplification', in R. Ordoubadian, and W. von Raffler Engel (eds.), *Views on Language*, Tennessee Inter-University Publishing, Murfreesboro, pp. 144–158.

Smolensky, Paul. 1996. *The initial state and 'Richness of the Base' in Optimality Theory* ms, Johns Hopkins University. [Available on Rutgers Optimality Archive, ROA-154.]

Smolensky, Paul, and Géraldine Legendre. 2006. *The Harmonic Mind: From Neural Computation to Optimality-Theoretic Grammar*, The MIT Press, Cambridge, MA.

Suzuki, K. 1998. *A Typological Investigation of Dissimilation*, unpublished Ph.D. dissertation, University of Arizona.

Tesar, Bruce and Paul Smolensky. 2000. *Learnability in Optimality Theory*. Cambridge, MA: MIT Press.

Treiman, Rebecca, Brett Kessler, Stephanie Knewasser, Ruth Tincoff, and Margo Bowman. 2000. 'English speakers' sensitivity to phonotactic patterns', in M. B. Broe and J. B. Pierrehumbert (eds.), *Papers in Laboratory Phonology V: Acquisition and the Lexicon*, Cambridge University Press, Cambridge, pp. 269-282.

Uhlenbeck, E.M. 1949. *De Structuur van het Javaanse Morpheem*, Nix, Bandung.

van den Berg, René. 1989. *A Grammar of the Muna Language*, Foris, Dordrecht.

van den Berg, René and La Ode Sidu. 1996. *Muna-English Dictionary*, KITLV Press, Leiden.

Yip, Moira. 1989. 'Feature Geometry and Co-occurrence Restrictions', *Phonology* **6**, 349-374.

Zuraw, Kie. 2000. *Patterned Exceptions in Phonology*, unpublished Ph.D. dissertation, UCLA.

Zuraw, Kie. 2002. 'Aggressive Reduplication', *Phonology* **19**, 395-439.

Appendix A: Features for similarity calculations

Dorsals

	k	g	$^{\text{ŋ}}$k	$^{\text{ŋ}}$g	ʁ	ŋ
sonorant	-	-	-	-	-	+
narrow	+	+	+	+	-	+
intermediate	-	-	-	-	+	-
wide	-	-	-	-	-	-
voice	-	+	-	+	+	+
nasal	-	-			-	+
prenasal	-	-	+	+		

Labials

	p	b	^{m}p	^{m}b	ɓ	f	w	m
sonorant	-	-	-	-	-	-	+	+
narrow	+	+	+	+	+	-	-	+
intermediate	-	-	-	-	-	+	-	-
wide	-	-	-	-	-	-	+	-
voice	-	+	-	+		-	+	+
nasal	-	-						+
prenasal	-	-	+	+				
implosive	-	-			+			
consonantal	+	+	+	+	+	+	-	+

Coronals

	t	d	ɗ	^{n}t	^{n}d	s	^{n}s	l	r	n
sonorant	-	-	-	-	-	-	-	+	+	+
narrow	+	+	+	+	+	-	-	-	-	-
intermediate	-	-	-	-	-	+	+	-	-	-
wide	-	-	-	-	-	-	-	+	+	+
voice	-	+		-	+	-	-	+	+	-
nasal	-	-								+
prenasal	-	-		+	+	-	+			
implosive	-	-	+							
lateral								+	-	

APPENDIX B: SIMILARITY VALUES FOR NON-IDENTICAL HOMORGANIC CONSONANTS OF MUNA

Dorsals

	k	**g**	**ᵑk**	**ᵑg**	**ʁ**
g	0.42				
ᵑk	0.56	0.31			
ᵑg	0.31	0.67	0.42		
ʁ	0.20	0.36	0.20	0.36	
ŋ	0.20	0.36	0.20	0.36	0.25

Labials

	p	**b**	**ᵐp**	**ᵐb**	**ɓ**	**f**	**w**
b	0.42						
ᵐp	0.60	0.31					
ᵐb	0.31	0.64	0.42				
ɓ	0.44	0.40	0.44	0.40			
f	0.33	0.18	0.33	0.18	0.29		
w	0.09	0.18	0.09	0.18	0.13	0.14	
m	0.17	0.36	0.17	0.36	0.22	0.11	0.43

Coronals

	t	**d**	**ɗ**	**ⁿt**	**ⁿd**	**s**	**ⁿs**	**l**	**r**
d	0.43								
ɗ	0.36	0.36							
ⁿt	0.43	0.25	0.36						
ⁿd	0.25	0.54	0.36	0.43					
s	0.42	0.21	0.20	0.21	0.13				
ⁿs	0.21	0.13	0.20	0.42	0.21	0.40			
l	0.07	0.15	0.11	0.07	0.15	0.09	0.09		
r	0.07	0.15	0.11	0.07	0.15	0.9	0.09	0.67	
n	0.14	0.33	0.22	0.14	0.33	0.08	0.08	0.38	0.38