Right-to-Left Biases for Vowel Harmony: Evidence from Artificial Grammar[*]

Sara Finley, William Badecker

Johns Hopkins University

1. Introduction

This paper uses artificial grammar learning data to examine the default nature of directionality in vowel harmony. We argue that vowel harmony is non-directional by default and that the right-to-left biases found within the typology of vowel harmony may best be thought of in terms of a bias against prefix harmony triggers.

Directionality has been an important issue within the theoretical discussion of vowel harmony because in many cases the same vowel harmony process can be described either directionally (e.g., from left-to-right) or non-directionally (e.g., stem-outward). For example, vowel harmony in Turkish appears to apply from left-to-right. However, because Turkish is a suffixing-only language, harmony can also be described in terms of a non-directional stem-outward harmony system. For most harmony languages, both interpretations may be descriptively adequate, but they make different predictions about what learners actually infer given ambiguous data. Therefore, the default nature of directionality could be either directional or non-directional (e.g., morphological). In the directional description, a vowel feature spreads in one direction from a designated source vowel (e.g., the rightmost vowel of the morphological stem) to a target vowel. In the non-directional description, the vowel feature spreads outward from the source vowel, regardless of the direction of spreading, so the vowel feature may spread from right-to-left, left-to-right, or both depending on the location of potential harmony undergoers.

The default nature of directionality (e.g., directional versus non-directional) in vowel harmony is still debated, as there are several empirical reasons to posit both directionality and non-directionality as the default. Directional descriptions are desirable for explaining the large number of languages that seem to only have harmony in one particular direction (e.g., Turkish, Finnish, Hungarian, Assamese). Within these

[*] We would like to acknowledge the following for vital discussion, comments and time on this project: Ariel Goldberg, Colin Wilson, Paul Smolensky, Bruce Hayes, Eric Bakovic, Luigi Burzio, John Kingston, Barbara Landau, Michael Wolmetz, the IGERT lab, and audiences at HOWL4, NELS38 and the 2008 LSA Annual Meeting.

languages, there appears to be an asymmetry between right-to-left harmony and left-to-right harmony, namely a bias towards right-to-left harmony (Hyman 2002). For example, languages such as Karajá (Riberio 2002), Assamese and Pulaar (Mahanta 2007) involve harmony from suffixes to stems (right-to-left), but not from prefixes to stems (left-to-right). Within the set of languages in which affix vowels may trigger harmony, there are cases where suffixes, but not prefixes spread to stems, (e.g., Turkana, Noske 2000), but no languages where prefixes, but not suffixes spread to stems (Bakovic 2000; Krämer 2002). In other words, prefix triggers imply suffix triggers. For example, Kalenjin has prefix triggers, but the majority of harmony triggers are suffix triggers (Hyman 2002). This lack of harmony from prefixes to stems suggests a directional description of harmony with a right-to-left bias. Further, related spreading processes such as consonant harmony (Hannson 2001) and tonal spreading (Chen 2000) are typically characterized in terms of specific directions.

However, there is also evidence in favor of a non-directional description of vowel harmony. First, it is possible to explain the default directionality of many harmony languages in terms of the morphology (Bakovic 2000), as in suffixing-only languages like Turkish and Hungarian. Secondly, there are many cases of dominant-recessive vowel harmony in which the presence of a particular feature value (e.g., [+ATR]) will induce spreading in both directions throughout the lexical item. These types of languages (e.g., Kalenjin, Nez Perce, Lango, Mayak) are best described non-directionally, because the dominant feature value spreads from stem to suffix and vice versa. It is the location of the feature value that drives spreading, rather than intrinsic directionality.

Because there is evidence in favor of both directional and non-directional descriptions of vowel harmony, other methods for exploring the default nature of directionality must be employed. To this end, we present data from two artificial grammar learning experiments that test for the default nature of directionality in vowel harmony. The artificial grammar learning paradigm is a promising approach for testing the nature of grammatical defaults. Many researchers (e.g., Finley and Badecker 2008, submitted; Moreton to appear; Pycha et al. 2003; Wilson 2003, 2006) have had success in training participants in novel languages, as well as testing particular details of linguistic theory. For example, the Poverty of the Stimulus Method (Wilson 2006) makes use of generalization to novel items in order to infer substantive learning strategies. Finley and Badecker (in preparation) trained participants on a stem-to-suffix height harmony rule with front vowel suffixes (e.g., [-mi]/[-me]; [bidi, bidimi]; [godo, godome]) or back vowel suffixes (e.g., [-mu]/[-mo]; [bidi bidimu]; [godo godomo]), and tested on both. They showed generalization to cross-linguistically preferred front vowel suffixes, but no generalization to dispreferred back vowel suffixes. Differential generalization of this sort provides evidence of learning biases, and can therefore shed light on the nature of default grammatical processes.

What makes the Poverty of the Stimulus Method so useful in testing learning biases is that the training data are designed to be ambiguous with regard to the generalization that they embody. For example, training with examples of height harmony with suffixes that only contain front vowels (e.g., [bidu bidu-gi], [bedo, bedo-ge]) is compatible with a harmony pattern that applies only to front vowels, as well as more general harmony pattern that applies to all vowels. The test data are disambiguating (contain front and back vowel suffixes) and thereby allow one to assess the nature of the generalization that the learners adopt (e.g., a harmony rule that applies to front vowels or a harmony rule that applies to both front and back vowels). This method allows the researcher to ask questions about learners' natural inferences on phonological processes. Here we use the Poverty of the Stimulus Method to test the default nature of directionality in vowel harmony. By exposing participants to prefixing-only or suffixing-only data during training, it is possible to test learners' generalization to novel affixes and therefore test generalization to a novel direction. If participants, exposed only to stems spreading to suffixes in a left-to-right harmony, infer a directional rule, they should not generalize to right-to-left harmony in novel prefixes. However, if the learner infers a non-directional rule (e.g., stem-outward), they should generalize to novel prefixes. By exposing the inferences that learners make when presented with a limited data set, the test data can inform us about the nature of directionality in vowel harmony.

2. Experimental Design

In Experiment 1, we expose participants to a stem-outward harmony training data with either prefixes only or suffixes only and then test generalization to the withheld (novel) affix type. Experiment 2 takes the same approach, but with affix-triggering harmony, where suffixes spread to stems or prefixes spread to stems. From this ambiguous training data, participants can either infer a strictly directional process, or a morphologically controlled non-directional process. Generalization to novel affix classes (e.g., prefixes if only trained on suffixes) suggests that the learner has inferred a non-directional harmony pattern.

For the experiments presented here, we ask two main questions. In Experiment 1, these questions are asked in reference to stem-outward harmony. In Experiment 2, these questions are asked in reference to affix-triggering harmony.

(1) What is the default nature of directionality in vowel harmony?

Regardless of the default nature of directionality in vowel harmony, it is possible for learners to prefer one direction over another. For example, an affix-driven harmony may prefer suffixes spreading to stems over stems spreading to suffixes. Because the typological data suggests a right-to-left preference in vowel harmony, we ask (2):

(2) Is there a right-to-left preference for vowel harmony?

These questions may be addressed through the behavior of our participants on items that target generalization of the harmony pattern to novel suffixes at test. If participants generalize vowel harmony both from suffixes to prefixes, and prefixes to suffixes, it suggests that they inferred a non-directional harmony pattern, and, crucially, that these participants are biased to learn non-directional harmony. If participants do not generalize to either novel prefixes or to novel suffixes, it suggests that they are biased to learn a strictly directional harmony pattern. If participants have a right-to-left bias for vowel harmony, learners should generalize asymmetrically; they should generalize to prefixes only in stem-controlled harmony, and to suffixes only in affix-triggering harmony.

Another measure of the right-to-left bias for harmony comes from the robustness of learning and generalization to prefixing and suffixing languages. For stem-controlled harmony, if participants show greater harmonic responses to prefixes than to suffixes, it suggests a right-to-left bias for harmony. For affix-triggering harmony, suggestion of a right-to-left bias is present if participants show greater learning for suffix harmony triggers (right-to-left) than prefix harmony triggers (left-to-right). Further, a right-to-left bias for harmony could be demonstrated if participants show greater generalization to novel stems than to either novel suffixes (in stem-controlled harmony) or novel prefixes (in prefix-triggering harmony).

In both experiments, participants in the critical conditions were exposed to a back/round harmony language. In this language, all lexical items contained either only front/unround vowels [i, e] or only back/round vowels [o, u]. All training items contained a harmonic two-syllable stem (stress on the initial syllable), followed by its concatenated form (prefixed or suffixed). Affixes were counterbalanced such that half of the participants heard an affix with a bilabial nasal at training ([mi]/[mu]) and a velar stop for novel affixes ([gi]/[gu]), and the other half tested on the opposite. For example, a participant trained on the prefix [mi-]/[mu-] would be tested on the suffix [-gi]/[-gu].

All participants were given a simple AXB perception task for English vowels. In this perception task, participants were given three isolated syllables. Participants were asked to indicate whether the second vowel was identical to the vowel in the first syllable (e.g., [bi pi ku]) or to the vowel in the third (e.g., [bi pu ku]). The task was used as a screening to ensure that all participants were able to discriminate isolated English vowels. If participants scored less than 70% on this task, the data from the test phase was discarded. No participants in Experiment 1 failed this task.

3. Experiment 1

Experiment 1 tests for the default nature of directionality of vowel harmony in stem-controlled harmony systems. We exposed participants to a back/round stem-controlled harmony pattern with either prefixes only or suffixes only. Test items included both prefixed and suffixed items in all conditions. Generalization to novel affix classes suggests a learning bias towards a non-directional harmony pattern.

3.1 Participants

48 adult native English speakers from the Johns Hopkins University community participated in Experiment 1. No participants had prior knowledge of a vowel harmony language, and none had participated in previous vowel harmony learning experiments. Participants were either paid $7 or given extra course credit for their participation.

3.2 Method

The experiment consisted of three phases: a training phase, a testing phase and the AXB perception phase, all run on a Macintosh computer using Psyscope (Cohen, MacWhinney, Flatt, and Provost 1993). Participants were given verbal instructions at the start of the experiment, and written instructions at the start of each phase. Before the training phase, participants were told that they would be listening to a language that they had never heard before, and that they should pay attention to how the language sounds, but need not memorize the words they hear. Participants were given no information about morphology or vowel harmony, nor were they given any feedback during the test.

There were three between-subject conditions in the training phase: Prefix Training, Suffix Training and Control conditions. Prefix Training participants were exposed to stems followed by prefixed items, Suffix Training participants were exposed to stems followed by suffixed items and Controls were exposed to a mixture of harmonic and disharmonic stems. The training phase for the critical training conditions consisted of 24 stem and stem+affix pairs repeated 5 times each in random order. Each bare stem (e.g., [bidi]) was followed by its affixed form with a 500 ms delay (e.g., [bidi, bidi-mi] for Suffix Training, [bidi, mi-bidi] for Prefix Training).

(3) Examples of Training Stimuli: Experiment 1[1]

Prefix Training	**Suffix Training**	**Control**
beme mi-beme	beme beme-mi	beme tipo
digi mi-digi	digi digi-mi	digi kino
gomo mu-gomo	gomo gomo-mu	gomo kote
nupu mu-nupu	nupu nupu-mu	nupu gome

The Control condition was designed to ensure that all results were due to learning rather than a particular bias in the stimuli or in English for a particular response independent of training. We exposed participants to both harmonic and disharmonic stems to give participants some exposure to the training set without providing the opportunity for these participants to make any inferences about a harmony (or disharmony) pattern in the data.

The test phase consisted of 36 forced-choice test item pairs. One item in the pair was harmonic, the other was disharmonic (e.g., [bidi-mi, *bidi-mu]). Items were counterbalanced for order of harmonic response and round versus unround affixes. There were three within-subject conditions in the test set that varied by training condition: Old Stems/Old Affix, New Stems/Old Affix and Old Stems/New Affix. The Old Stems/Old Affix condition contained items that were identical to the training set. The New Stems/Old Affix condition contained items that had the same affix as in the training set, but a novel stem. These items were used to test for abstract learning of the directional

[1] Note that the same stems appeared in the training and test conditions for participants who heard [gu]/[gi] affix alternations at training.

harmony pattern. The Old Stems/New Affix condition contained stems that were heard in the training set, but affixes that were held-out. This was to test for generalization to novel affixes and the default nature of directionality in vowel harmony. Half of the participants in the Control condition heard test items identical to those in the Suffix Training condition, and the other half heard test items identical to those in the Prefix Training condition. The task concluded with the AXB perception task described above.

(4) Examples of Test Stimuli: Experiment 1[2]

	Old Stems/Old Affix		**New Stems/Old Affix**		**Old Stems/New Affix**	
Prefix Training	mi-beme mi-digi mu-gomo	*mu-beme *mu-digi *mi-gomo	mi-tede mu-bugu mu-pogo	*mu-tede *mi-bugu *mi-pogo	beme-gi digi-gi gomo-gu	*beme-gu *digi-gu *gomo-gi
Suffix Training	beme-mi digi-mi gomo-mu	*beme-mu *digi-mu *gomo-mu	tede-mi bugu-mu pogo-mu	*tede-mu *bugu-mi *pogo-mi	gi-beme gi-digi gu-gomo	*gu-beme *gu-digi *gi-gomo

3.3 Results

Results of Experiment 1 are reported in the figure in (5) below in terms of mean proportion harmonic responses. Each test condition is reported separately. To test the overall effect of training, we compared each training condition to the Control condition separately with a 2 x 3 mixed design ANOVA. The ANOVA comparing the Suffix Training condition to the Control condition showed a significant effect of Training ($F(1, 30) = 29.11$; $p < 0.001$), but no significant effect of Test Item ($F(2, 60) = 1.03$; n.s.). The results show that there was a significant effect of Training on the Suffix Training condition; participants learned the harmony pattern. The ANOVA comparing the Prefix Training condition to the Control showed a significant effect of Training ($F(1,30) = 18.36$; $p < 0.001$), but no effect of Test Item ($F(2,60) = 1.34$; n.s.). These results suggest that overall, participants learned the harmony pattern.

To assess the level of generalization to novel affixes, we compared the means of the Old Stem/New Affix Test Condition and the Control in each of the Training Conditions. There was a reliable difference for both novel suffixes (the New Affix condition for the Prefix Training Condition) ($t(30)=2.10$; $p < 0.05$) and novel prefixes (the New Affix condition for the Suffix Training Condition) ($t(30)=2.10$; $p < 0.05$). In both of these conditions, the mean harmonic responses differed significantly from Control participants. This suggests that learners inferred a non-directional, morphologically controlled (e.g., stem-outward) rule; both training conditions generalized to novel affixes.

To assess the robustness of generalization for novel affixes, we compared the means of the New Affix and New Stem test conditions for the Prefix and Suffix Training

[2] All items were counterbalanced such that harmonic items appeared first in the forced choice test half of the time, and front vowel items appeared first half of the time.

Conditions combined. There was a significant effect of Test Item ($F(1,30)=12.20$; $p < 0.01$), suggesting that generalization to novel stems was more robust than generalization to novel affixes. There was no effect of Training ($F < 1$), and no interaction between the New Stem and the New Affix conditions ($F < 1$), suggesting no preference for novel prefixes over novel suffixes. These results demonstrate an effect of directionality, in that generalization to a novel direction is less robust than generalization to novel stems.

Our results show significant effects of Training over all three conditions of test items, as well as the New Affix condition for both Prefix Training and Suffix Training conditions. This generalization to novel affixes for novel prefixes as well as for novel suffixes suggests that participants inferred a non-directional harmony pattern.

(5) Experiment 1 Results

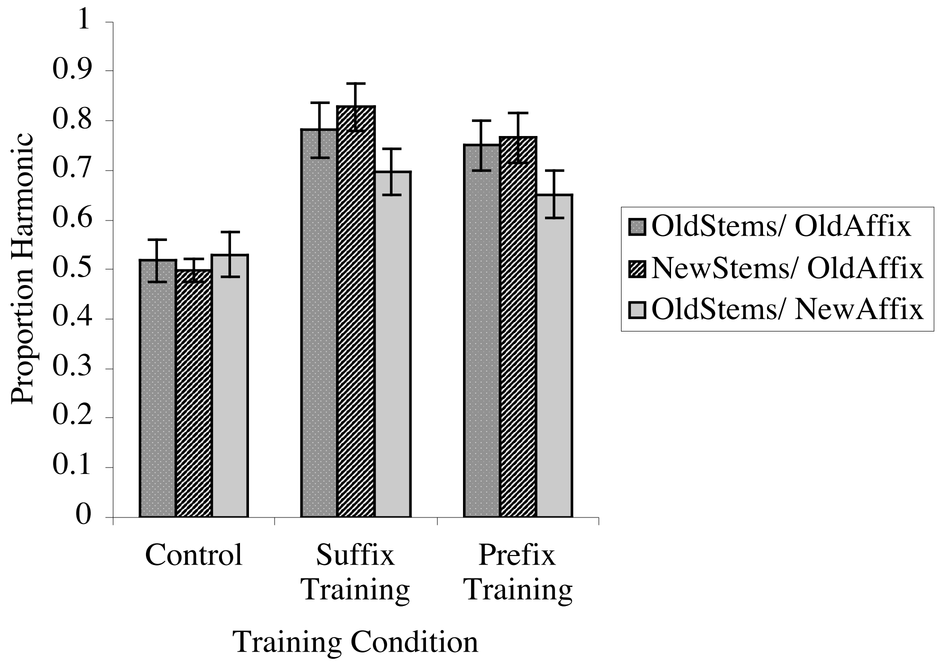

3.4 Discussion

The results of Experiment 1 provide evidence for non-directionality in vowel harmony, in that participants were able to generalize to novel affixes. However, participants were more likely to prefer harmonic items for novel stems than harmonic items in a novel direction. This may be because learners infer the harmony rule based on the morphological alternations (e.g., [-mi]/[-mu]) that they are exposed to, and thus perceive novel stems with old affixes as closer to the exemplar used for training than old stems with a novel affix. This finding is consistent with previous results in which participants were less likely to choose the harmonic response for items with a novel suffix vowel than items with a novel stem vowel (but the same suffix vowel) (Finley and Badecker 2008).

Generalizing to a novel affix, especially a novel class of affixes (prefix to suffix and vice versa) is more difficult than generalizing to a novel stem.

One possible interpretation of these results is that the selection of harmonic affixes in the New Affix conditions actually reflects a harmony rule with absolute directionality in which the novel affixes trigger spreading to a novel stem. For example, if a learner selects the harmonic item [mi-bidi], it is because they are following a left-to-right harmony rule in which the initial vowel spreads its features to the stem. If participants learned a directional rule in which affixes could spread to stems, they may still select the harmonic option in the forced-choice test. However, the design of the experiment and test items provides little reason to believe that this is the strategy that learners took. First, participants were only exposed to harmonic alternations in which an affix alternated between two different allomorphs. Thus, there is little reason to suspect that learners inferred a harmony pattern in which affixes spread to stems. Second, all test items differed only in terms of the vowel quality of the affix (e.g., [mi]/[mu]). Therefore, it is implied in the test items that it is the affix that undergoes harmony. Further, because participants were familiarized with all stems in the New Affix condition, it is unlikely that learners would infer that the affix vowel was the harmony trigger.

Our results provided little support for a right-to-left bias for stem-outward harmony. First, participants generalized to both prefixes and suffixes, and there was no significant difference between the rate of generalization to prefixes and suffixes. Typologically, this result is not surprising. The right-to-left biases for vowel harmony are seen in affix-dominant harmony, not stem-controlled harmony. For example, languages that have right-to-left harmony only (e.g., Assamese, Pulaar) are all affix-dominant harmony systems. Second, the right-to-left bias found in harmony may be due to a bias against prefix triggers for harmony, as the only known cases of languages where prefixes spread harmony also allow suffixes to spread harmony. Therefore, a true test of the right-to-left bias for vowel harmony would be found using affix-triggering harmony.

4. Experiment 2

Experiment 2 uses the same methodology as Experiment 1, but with an affix-triggering harmony language, to assess the default nature of directionality in vowel harmony, and the right-to-left bias for harmony found cross-linguistically.

4.1 Participants

38 adult native speakers of English participated in Experiment 2 for extra course credit. No participant participated in Experiment 1 or a previous harmony experiment. Two participants were dropped due to a failure meet threshold (70%) on the AXB perception task, for a total of 36 participants used in the data analysis (12 participants in each training condition).

4.2 Method

The procedure for Experiment 2 was identical to Experiment 1, with some minor changes reflecting the fact that Experiment 2 involved training participants on an affix triggered harmony pattern. Suffixes/prefixes spread to stems, rather than spreading from the stem to an affix. In order to present participants with alternations, all stems were unround in their bare form, and all affixes were round. Thus, participants heard a bare unround stem followed by its concatenated round form (e.g., [bidi, mu-budu], [bidi, budu-mu]).

(6) Examples of Training Stimuli: Experiment 2

Prefix Training	**Suffix Training**	**Control**
beme mu-bomo	beme bomo-mu	beme bumi
kene mu-kono	kene kono-mu	kene kino
midi mu-mudu	midi mudu-mu	midi nego
pidi mu-pudu	pidi pudu-mu	pidi podi

As in Experiment 1, there were three training conditions: Prefix Training, Suffix Training and Control. The Prefix Training condition heard prefix harmony triggers (e.g., [bidi, mubudu]). Participants in the Suffix Training condition heard suffix harmony triggers (e.g., [bidi budumu]). The Control condition heard a mixture of harmonic and disharmonic stems (e.g., [bidi, pegu]). The forced choice test items reflected the fact that alternations took place in the stem rather than the affix. Thus, the affix always contained a round vowel, but the stem alternated between round and unround (e.g., [bidi-mu, budu-mu], [mu-bidi, mu-budu]). Test items were of the same three conditions as Experiment 1: Old Stems/Old Affix (which were items that appeared in the training set), New Stems/ Old Affix (which were items that had the same affix as training but novel stems) and Old Stem/New Affix (which were items that contained a novel affix) and were used to test for generalization to a novel affix.

(7) Examples of Test Stimuli: Experiment 2

Prefix Training	**Suffix Training**
Old Stems/Old Affix	**Old Stems/Old Affix**
*kene-mu kono-mu	*gu-kene gu-kono
*pidi-mu pudu-mu	*gu-pidi gu-pudu
New Stems/Old Affix	**New Stems/Old Affix**
dono-mu *dene-mu	gu-dono *gu-dene
bugu-mu *bigi-mu	gu-bugu *gu-bigi
Old Stems/New Affix	**Old Stem/New Affix**
*gu-kene gu-kono	*kene-mu kono-mu
*gu-pidi gu-pudu	*pidi-mu pudu-mu

The experiment concluded with the same AXB perception task as Experiment 1, used as a screening to ensure that all could perceive the training and test stimuli. Two participants failed this task, and their learning data was discarded.

4.3 Results

Results of Experiment 2 are reported in the figure in (8) below in terms of mean proportion harmonic responses. Each test condition is reported separately. To test the overall effect of training, we compared each training condition to the Control condition separately with a 2 x 3 mixed design ANOVA. The ANOVA comparing the Suffix Training condition to the Control condition showed a significant effect of Training ($F(1, 22) = 10.2$; $p < 0.05$), but no effect of Test Item ($F(2,44) = 1.03$; n.s.). These results indicate that participants successfully learned the harmony pattern. The ANOVA comparing the Prefix Training condition to the Control showed no effect of Training ($F<1$), no significant effect of Test item, and a significant interaction between factors ($F(2,44) = 5.09$; $p < 0.05$). This indicates that participants did not learn the harmony pattern in the Suffix Training condition.

(8) Experiment 2 Results

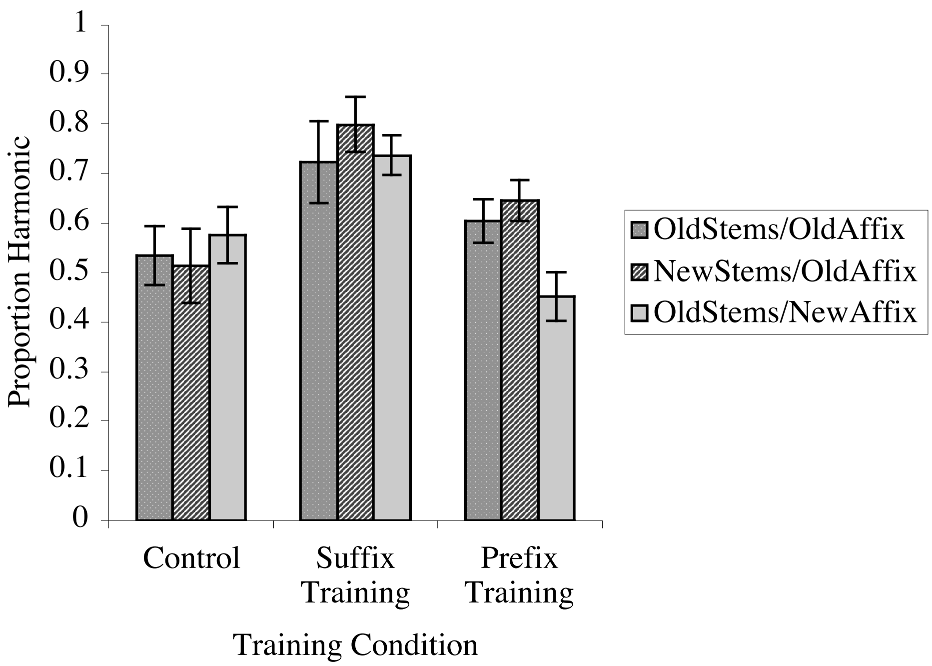

To assess the level of generalization to novel affixes, we compared the means of the Old Stem/New Affix Test condition to the Control conditions in each of the Training conditions. Novel prefixes (the New Affix condition for the Prefix Training Condition) were significantly different from the Control Condition ($t(22)=2.10$; $p < 0.05$). There was a marginal effect of Training in the Novel Affix condition of the Prefix Training condition, but in the opposite direction. There were slightly fewer harmonic responses to novel affixes (suffixes) than controls ($t(22)=2.0$; $p = 0.06$). These results suggest that if one is able to learn a harmony rule, then that harmony rule will be non-directional by default.

We compared the Prefix Training condition to the Suffix condition to test whether or not learning or generalization was more robust in either condition, using the same mixed-design ANOVA as above. There was a significant effect of Training ($F(1,22) = 6.4$; $p < 0.05$) and no significant interaction ($F(2, 44) = 2.0$; $p > 0.05$). These results suggest that learning was more robust for the Suffix Training condition than the Prefix Training Condition. There was a significant effect of Test Item ($F(2,44) = 4.30$; $p < 0.05$); there was an overall significant difference between New Stems and New Affix ($F(2, 44) = 10.37$ $p < 0.01$). This result supports a right-to-left bias for vowel harmony in the sense that participants seemed biased against learning a harmony pattern with only prefix triggers.

In order to ensure that the null effect of training in the Prefix Training condition was not due to the fact that there were fewer harmony responses for New Affix, we compared the Prefix Training condition with the Control condition for just New Stems and Old Stems/Old Affix test conditions. There was no main effect for this ANOVA ($F(1, 22) = 1.89$; $p > 0.05$), no effect of Test item ($F < 1$), and no interaction ($F < 1$). We also compared the Prefix Training condition to 50% chance. There was also no significant effect of Training ($F(1, 22) = 2.01$; $p > 0.05$). These results suggest that participants in the Prefix Training condition did not learn the harmony pattern.

4.4 Discussion

The results of Experiment 2 are consistent with the results of Experiment 1: participants generalized to novel affixes if they learned the harmony pattern. Participants in Experiment 2 were unable to learn the harmony pattern with prefix harmony triggers. This indicates that our learners were biased against prefix-triggering harmony, notably, this bias is found cross-linguistically: the few known languages that allow prefix harmony triggers predominantly allow suffix harmony triggers. This is consistent with our findings that learners only generalize to novel prefix triggers from suffix harmony triggers.

5. General Discussion

The experiments presented in this paper addressed two questions concerning the representation of vowel harmony: is vowel harmony directional by default? And is there a right-to-left bias for vowel harmony? The results of our experiments support a non-directional default for harmony, and a right-to-left bias for directional harmony. However, this right-to-left bias manifests itself as a bias against prefix harmony triggers.

Stem controlled harmony is non-directional by default, and there is little evidence for a right-to-left bias for stem-controlled harmony. Affix-triggering harmony appears to be more complex. Suffix-triggering harmony is learnable, while prefix-triggering harmony is not learnable within the short training period that we provided for our participants. However, suffix-triggering harmony does seem to imply a general, non-directional affix-triggering harmony pattern, as learners exposed to suffix triggers also accepted prefix harmony triggers. Thus, it appears that prefix triggers are acceptable only in the presence of suffix harmony triggers, a finding consistent with cross-linguistic patterns. However, training on prefix harmony triggers is not enough to induce acceptability of suffix harmony triggers. This right-to-left bias in affix triggering harmony may arise from co-articulatory pressures (Ohala 1994). Vowel-to-vowel co-articulation is typically from right-to-left. Further, there is evidence that in some modes of speech, speech errors are most often regressive (right-to-left) (Fromkin 1973).

One potential cause for the difficulty that participants had in learning the prefix trigger harmony pattern is the fact that the stress was on the first syllable in the stimulus items. While all vowels were clearly articulated and had at least secondary stress (to avoid vowel reduction), the main stress was always on the initial syllable. For prefixed items, the primary stress shifted from the stem to the prefix, which may have made parsing the prefixed item more difficult than parsing the suffixed item. However, if learning failure arose from parsing errors, then we should expect (1) difficulty in learning for stem-controlled prefixed forms and (2) no generalization to novel prefix harmony triggers. However, we found both robust learning of prefixed forms for stem-controlled harmony as well as robust generalization to novel prefixed forms. Thus, there must be something that made the prefix triggering harmony pattern difficult to learn beyond the stress placement.

5.1 Implications for a Theory of Vowel Harmony

The experimental evidence that vowel harmony is non-directional by default has important consequences for a theory of vowel harmony. Within Optimality Theory (Prince and Smolensky 1993/2004), our results support a set of constraints and representations that predict a typology of vowel harmony languages that show a bias for non-directionality and against prefix harmony triggers.

One possibility for predicting a bias against prefix harmony triggers may come from a constraint against prefixes as the source of harmony. Within Headed-Feature Domains Theory (Smolensky 2006) and Span Theory (McCarthy 2004), the source for spreading is represented as a head in the spreading domain. Because prefixes tend to be phonologically weak in that they are more likely to undergo processes than to trigger them, a constraint that penalizes spreading domains in which a prefix is the source or head is a worthwhile option to pursue further.

There are several possibilities for deriving the default non-directional nature of vowel harmony. One approach is to employ a non-directional constraint such as AGREE (Bakovic 2000), where harmony is induced by agreement rather than a particular source of spreading. A similar option is to employ a non-directional harmony constraint such as those found in Span Theory and Headed-Feature Domains Theory, but allow for the possibility of specific direction through a set of additional (alignment) constraints and therefore induce a preferred direction of spreading that allow for directionality. A further option is to use inherently directional constraints such as Align-R/L or Spread-R/L (e.g., Kirchner 1993), but with a restriction on the initial state in learning that would induce a non-directional harmony when training data is ambiguous between directional and bi-directional vowel harmony. For example, if Spread-R and Spread-L were, at the initial state, paired together with no ranking between the two constraints, then as faithfulness moves to a lower-ranked position, it will move below both Spread-R and Spread-L unless there is evidence of a specific ranking between the two constraints. One advantage of this approach is that it directly encodes the results of the experiment as a learning bias, leaving open the possibility for languages with a specific direction for spreading. However, more research is needed to differentiate between these different formalizations of the non-directional bias for vowel harmony.

6. Conclusion

This paper provides evidence that artificial grammar learning experiments can shed light onto the nature of phonological knowledge, specifically the learning biases that lead to patterns in phonological typologies. The experiments presented in this paper support the notion of a right-to-left bias for vowel harmony as a bias against prefix harmony triggers, as well as the non-directional nature of vowel harmony.

This work also represents the benefits of an interaction of experimental and theoretical methodologies. Theoretical considerations created the need for the experiments in this paper. The results of our experiments create new issues for theoretical phonology: how to represent the default nature of directionality in vowel harmony, and the bias against prefix harmony triggers. Future work will continue in this vein, continuously integrating theoretical and experimental methodologies to better understand the nature of phonological representations.

References

Bakovic, Eric. 2000. Harmony, dominance and control. Doctoral dissertation, Rutgers University.

Chen, Matthew. 2000. Tone sandhi: Patterns across Chinese dialects. Cambridge: Cambridge University Press.

Cohen, J. D., MacWhinney, B., Flatt, M., and Provost, J. 1993. PsyScope: A new graphic interactive environment for designing psychology experiments. Behavioral Research Methods, Instruments and Computers 25, 257-271.

Finley, Sara, and Badecker, William. 2008. Substantive biases for vowel harmony. To appear in Proceedings of BLS 33.

Finley, Sara, and Badecker, William. submitted. Artificial grammar learning and feature-based generalization. Ms, Johns Hopkins University.

Finley, Sara, and Badecker, William. In preparation. Front/back asymmetries in height harmony. Ms, Johns Hopkins University.

Fromkin, Victoria. 1973. Speech errors as linguistic evidence. Berlin: Walter de Gruyter.

Hannson, Gunnar O. 2001. Theoretical and typological issues in consonant harmony. Doctoral dissertation, University of California, Berkeley.

Hyman, Larry H. 2002. Is there a right-to-left bias in vowel harmony? In International Phonology Meeting 9.
Krämer, Martin. 2002. The root-affix metaconstraint and its reversal in vowel harmony. Paper presented at the 10th Manchester Phonology Meeting, May 2002.
Mahanta, Shakuntala. 2007. Directionality and locality in vowel harmony: With special reference to vowel harmony in Assamese. Doctoral dissertation. Utrecht Institute of Linguistics.
McCarthy, John. 2004. Headed spans and autosegmental spreading. Ms, Umass, Amherst.
Moreton, Elliot. to appear. Analytic bias and phonological typology. Phonology.
Nokse, M. 2000. [ATR] harmony in Turkana: A case of suffix faith >> root faith. NNLT 18, 771-812.
Ohala, John. 1994. Hierarchies of environments for sound variation; plus implications for 'neutral' vowels in vowel harmony. Acta Linguistica Hafniensia 27, 371-382.
Prince, Alan, and Smolesnsky, Paul. 1993/2004. Optimality Theory: Constraint interaction in generative grammar. Cambridge, Mass.: Blackwell.
Pycha, Anne, Nowak, E., Shin, and Shosted, Ryan. 2003. Phonological rule-learning and its implications for a theory of vowel harmony. In Proceedings of WCCFL 22, 101-113.
Riberio, E. R.2002. Directionality in vowel harmony: The case of Karajá (Macro Je). In Proceedings of BLS 28.
Smolensky, Paul. 2006. Optimality in phonology II: Harmonic completeness, local constraint conjunction and feature-domain markedness. In The harmonic mind: From neural computation to Optimality-Theoretic grammar, Volume II., eds. Paul Smolensky and Geraldine Legendre, 27-160. Cambridge, Mass: MIT Press.
Wilson, Colin. 2003. Experimental investigations of phonological naturalness. In Proceedings of WCCFL 22, 101-114.
Wilson, Colin. 2006. Learning phonology with substantive bias: An experimental and computational study of velar palatalization. Cognitive Science 30, 945-982.

Sara Finley
Department of Cognitive Science
3400 N. Charles St.
Johns Hopkins University
Baltimore, MD 21218

finley@cogsci.jhu.edu

William Badecker
Department of Cognitive Science
3400 N. Charles St.
Johns Hopkins University
Baltimore, MD 21218

badecker@ jhu.edu