

Directionality Effects and Exceptions in Learning Phonological Alternations

Anqi Wang and Sara Finley*

1 Introduction

An important question in the cognitive science of language is whether the typological tendencies found in the world's languages are shaped by learning biases, whereby patterns that are cross-linguistically common are easier to learn than patterns that are rare or unattested. A promising methodology to explore how learning biases might shape linguistic typology is the artificial language learning paradigm, in which naïve learners are exposed to a miniature version of a language in a short period of time and tested on their learning and generalization of that language. Recent research has suggested a possible link between the typology of cross-linguistic patterns and learnability (e.g., Culbertson et al. 2012; Finley 2012). In phonology, a large body of research relating typology to learnability has focused on vowel harmony, a phonological pattern found in many languages of the world, in which adjacent vowels in a lexical item agree in some phonological feature (van der Hulst 2016), such as rounding or backness. Previous research has shown that adult participants are able to acquire a back/round vowel harmony pattern with relatively minimal exposure to the patterns (Finley to appear) and may be biased to learn vowel harmony patterns that are both phonetically natural and typologically frequent (Finley 2012).

The typology of vowel harmony systems suggests an interaction between directionality and morphology. While phonetic factors favor right-to-left spreading (Hyman 2002), morphological factors favor spreading from a stem to a suffix (left-to-right) (Bakovic 2000). In addition, vowel harmony is rife with exceptions, such that many affixes fail to participate in vowel harmony (Finley 2010). While previous research has explored learning biases for directionality in vowel harmony (e.g., Finley and Badecker 2009, White et al. 2018) and exceptionality in vowel harmony (e.g., Baer-Henney et al. 2015, Finley to appear), the current study aims to explore the interaction of directionality and exceptions in vowel harmony.

This paper is organized as follows: Section 2 reviews the background of the current study and Sections 3 and 4 present the two experiments we conducted to test our hypotheses. We then provide a general discussion of the results in Section 5. Section 6 concludes.

2 Background

Most phonological alternations display some degree of exceptionality (Coetzee and Pater 2011, Finley 2010), resulting in failure to reduce some language data to one general pattern. Those irregularities in language data pose challenges for generative linguistics, as explaining exceptions makes those models more complicated (Haut 2020). Language learning can be more challenging with those models, as individuals need to predict which instances can be inferred from the regular, consistent phonological pattern and which cannot.

*We would like to thank Ariel Goldberg for lending his voice, our participants, as well as our audiences at PLC, WPA, and CogSci 2021. We are also grateful to the Severtson Foundation Fellowship from PLU, which funded this research. All errors are our own.

The present study concentrates on vowel harmony, a phonological pattern found in many languages of the world (but not English), in which, all vowels in a word agree with each other about a specific phonetically based feature, such as front/back or high/low (van der Hulst 2016). For instance, in a back/round vowel harmony system, if the first vowel in a word is front (e.g., /i/, /e/), all the other vowels in this word must also be front. Similarly, words that begin with a back/round vowel (e.g., /o/, /u/) must only contain back/round vowels. For example, in a simplified account of Hungarian, the dative suffix has two versions (allomorphs), /-nək/ or /-nək/. When the stem vowel is front, the suffix vowel is also front, as in [tsi:m-nək] ‘address-DAT’, and when the stem vowel is back, the suffix vowel is also back, as in [ɔ̀bɫək-nək] ‘window-DAT’ (Hayes and Londe 2006).

Like Hungarian, many cases of vowel harmony are manifested as morphophonological alternations. These morphological alternations can also be described in terms of directionality. Vowel harmony languages display both left-to-right and right-to-left characteristics. However, directional spreading can often be explained in terms of the morphological properties of the language (Bakovic 2000). For instance, vowel harmony in Turkish, a suffixing-only language, seems to apply rightward from a stem to a suffix.

Because there is a cross-linguistic bias towards suffixes over prefixes (Cutler, Hawkins, and Gilligan 1985), many languages show left-to-right vowel harmony, and there is some evidence of a bias for harmony to apply to suffixes over prefixes (White et al. 2018). However, phonetic factors may also favor right-to-left harmony (Hyman 2002), and learners may be biased towards right-to-left harmony when spreading applies from affix to stem rather than from stem to affix (Finley and Badecker 2009). In addition, learners’ biases towards directional spreading patterns over majority rules hold for both linguistic and non-linguistic stimuli, suggesting that these constraints might influence the distribution of linguistic patterns across the world’s natural languages (Finley and Badecker 2010).

Because most phonological alternations show some aspects of exceptionality, it is crucial to understand how exceptionality affects language learning. Previous studies exploring the learnability of exceptionality in phonological patterns (e.g., Coetzee, 2009) have suggested a bias towards non-alternation over alternation. For instance, Coetzee (2009) trained participants on singular-plural alternations, in which one plural form did not change, while the final vowel in the counterpart became tense in 75% of the items in the training set. Participants were able to extend the non-alternating pattern to novel items when trained on data with equal evidence for alternation and non-alternation. Additionally, learners displayed a bias towards alternation when the frequency structure of learning data presented changed. This indicates that learners are inclined to extend non-alternation to novel items, and lexical items that alternate must be marked as such (Coetzee 2009).

However, when exposed to exceptions in vowel harmony learners may be biased towards harmony over non-alternation (Finley to appear). And this requires many instances of harmony over disharmony to learn the harmonic patterns (Baer-Henney et al. 2015). One question is whether learners might make inferences about the directionality of a harmony pattern when exposed to instances that contained both alternating and non-alternating affixes. The current study thus explores the interaction of directionality and exceptions in vowel harmony using an artificial language learning paradigm.

The present study makes use of a similar paradigm to Finley and Badecker (2009) and Finley (to appear) by exposing American-English speaking participants to a back/round harmony pattern that includes both alternating and non-alternating affixes, and tests for generalization to the novel direction (e.g., from prefixes to suffixes and from suffixes to prefixes). In Experiment 1, participants were exposed to a language with either prefixes or suffixes, where one affix alternated and the other did not, and then tested on the novel direction. If participants are biased towards suffixes as harmony undergoers, they should be more likely to generalize novel affixes from prefixes to suffixes than the opposite direction. If participants are biased towards right-to-left

harmony, they should be more likely to generalize from suffixes to prefixes.

In Experiment 2, participants were exposed to a harmony language that is ambiguous between a directional pattern (e.g., stem to suffix harmony) and a pattern with exceptions (e.g., [me]/[mo] alternates while [go] fails to alternate). If participants exhibit a bias towards suffixes undergoing vowel harmony (White et al. 2018), participants should treat the non-alternating prefixes as an indicator that harmony does not apply to prefixes and fail to extend the alternating affix to prefixes. If there is a right-to-left bias for harmony (Hyman 2002), participants should be able to treat the non-alternating suffixes as an indicator that harmony does not apply to suffixes and fail to extend the alternating prefixes to suffixes.

3 Experiment 1

The present study utilized an artificial language learning paradigm in which learners were trained on a novel back/round vowel harmony pattern with two affixes, one that alternated based on vowel harmony, and another that did not alternate, regardless of the feature value of the stem vowels. As in Finley (to appear), nonce words were presented as a repeated series of triads with a stem (in isolation) followed by two affixed forms. Specifically, one affix (either prefix or suffix, depending on the condition) alternated between [me] and [mo] depending on the back/round features of the stem, and the other affix was [go] regardless of the phonetic features of the stem. Following the exposure phase, learning and generalization was assessed using a two-alternative forced choice test, specified below.

3.1 Participants

All 40 participants were adults who were fluent in English speakers recruited from the psychology department’s research participation pool at a private institution in the Pacific Northwest of the USA. No participants had any previous experience with a vowel harmony system, natural or artificial. Participants were randomly assigned to one of two training conditions: PrefixOnly and SuffixOnly.

PrefixOnly	SuffixOnly
bono mobono gobono	bono bonomo bonogo
deke medeke godeke	deke dekeme dekego
nepe menepe gonepe	nepe nepeme nepego

Table 1: Examples of Training Items in Experiment 1.

3.2 Design and Materials

Participants were exposed to a miniature language with a back/round vowel harmony pattern with exceptions. The pattern was presented in a set of three words (triads): stem, stem+affix_{alternating}, stem+affix_{non-alternating} (e.g., [bono, mobono, gobono] in the PrefixOnly condition). Stems triggered an alternating affix (either a prefix or a suffix, depending on the condition) that was either [me] or [mo] depending on whether the vowels in the stems, which included front vowels (/i/ or /e/) or back vowels (/o/ or /u/). In the case of the non-alternating affix, the affix was always [go] regardless of whether the stem vowels were front or back. All stems took the form CVCV (e.g., [nepe]) with the vowels following back/round vowel harmony constraints, and the consonants drawn from the set (/p, t, k, b, d, g, m, n/). See Table 1 for examples of training stimuli.

All stem vowels were either front or back and were always harmonic. There were 24 triads, repeated five times each in a random order. In the PrefixOnly condition, both alternating affixes were prefixes (e.g., [bide, mebide, gobide]); in the SuffixOnly

condition, both alternating and non-alternating affixes were always suffixes (e.g., [bide, bideme, bidego]). These two conditions were ambiguous between directionality (e.g., stem to suffix harmony) and exceptionality (e.g., [me/mo] alternates, while [go] fails to alternate).

Learning was assessed via a 48-item two-alternative, forced-choice test, with 4 test conditions, and 12 items in each condition (see Table 2 for examples of test items). Stems were identical for both choices, except for the final vowel: either /e/ or /o/. Because all the stems were harmonic, the choice of the last vowel (/e/ or /o/) relied on the back/round features of the vowels in the first two syllables. All test items contained novel stems that were not heard in the training stage. All the participants responded to the same items, regardless of training condition.

New_Me suffix items and Me_New prefix items contained stems not heard in the training set, with the [me/mo] suffixes and prefixes. These items were designed to test for the learnability and generalization of the alternating affixes. New_Go suffix items and Go_New prefix items contained novel stems with the [go/ge] suffixes and prefixes. These items were designed to test for the learnability and generalization of the non-alternating affixes.

Test Item	Front Vowel Affix vs. Back Vowel Affix
Me_New prefix	metede vs. *motede *mekopo vs. mokopo
Go_New prefix	*getede vs. gotede *gekopo vs. gokopo
New_Me suffix	tedeme vs. *tedemo *kopome vs. kopomo
New_Go suffix	*tedege vs. tedego *kopoge vs. kopogo

Table 2: Examples of test items; * Indicates the incorrect/ungrammatical response.

All stimuli were recorded by an adult male speaker of American English with control for syllable structure. Stress was always on the final syllable using English pronunciations. All stimulus items were normalized for intensity (set at 70dB). All stimuli creation and modification to sound files was performed in Praat.

3.3 Procedure

All phases of the experiment were presented online using a web-based presentation platform (FindingFive Team 2019). Participants were asked to listen to the words from a miniature, made-up language using headphones in a quiet location. On each training trial, the participants heard the three words in the triad, and were required to click ‘Continue’ to move to the next trial.

Immediately following the training phase, participants were given a two-alternative, forced-choice test, with the test items described above presented in a random order. Participants were asked to select the word (first or second) that most likely belonged to the language. Participants were instructed to wait until both items played before making their selection. Upon completion of the experiment, all participants were debriefed; the entire experiment took approximately 20 minutes to complete.

3.4 Results

Mean proportion of correct responses for all test items, and the standard errors of the mean (derived from the function `summarySE` in R), are given in Figure 1. All trials that were longer than 10s were removed for analysis. The data were fitted into mixed effects

logistic regression models with Laplace approximation using the lme4 package (Bates et al. 2015) in R Studio. Pairwise comparisons were using the emmeans/lsmmeans packages (Lenth 2016, 2018). Due to convergence errors, the models generally only included random intercepts for Subjects and Items, unless otherwise noted.

Participants in the PrefixOnly condition and SuffixOnly condition were more likely to select the correct ([go]) forms for both prefixes (mean = 0.72 ± 0.041) and suffixes (mean = 0.73 ± 0.041), than the correct ([me/mo]) forms for alternating prefixes (mean = 0.62 ± 0.045) and suffixes (mean = 0.62 ± 0.045) ($\beta = 0.62$, $SE = 0.27$, $z = 2.30$, $p = 0.021$), based on a model with [me/mo] items set as the baseline, with random intercepts for Subjects and Item, and random slopes for Subjects by test item. This model showed a significant intercept ($\beta = 0.62$, $SE = 0.19$, $z = 3.34$, $p < 0.001$), suggesting that participants learned the behavior of the alternating affixes, even though there were significantly more correct responses to non-alternating affixes.




Figure 1: Experiment 1 results (means and confidence intervals of the mean).

To test if participants exhibit directionality biases, and if learners were able to generalize to novel directions, separate Training by Test condition interaction models were created with different baselines: $\text{glmer}(\text{response_correct} == 1 \sim \text{Test} * \text{Training} + (1|\text{stimuli_presented}) + (1|\text{participant_id}), \text{family} = \text{"binomial"}, \text{data} = \text{Pre_Suffix})$. A significant intercept was interpreted as significantly different from 50% chance. Due to space constraints, only theoretically relevant comparisons are included. However, full data and analysis code can be found at: <https://osf.io/upx5m/>.

The model with New_Me_Suffix items and the PrefixOnly condition as the baseline did not show a significant intercept, suggesting that participants in the PrefixOnly condition did not generalize the alternating affixes from prefixes to suffixes ($\beta = 0.31$, $SE = 0.23$, $z = 1.31$, $p = 0.19$). The comparison of New_Go_Suffix items and New_Me_Suffix items revealed that there were significantly more correct responses to New_Go_Suffix items (mean = 0.69 ± 0.063) compared to New_Me_Suffix items (mean = 0.57 ± 0.068) ($\beta = 0.57$, $SE = 0.23$, $z = 2.44$, $p = 0.015$), indicating that participants were more likely to make generalizations to the novel direction (new suffix) for non-alternating [go] items than alternating [me/mo] items.

The model with New_Me_Prefix items and the SuffixOnly condition as the baseline showed a significant intercept, suggesting that participants in the Suffix Only condition generalized the alternating affixes from suffixes to prefixes ($\beta = 0.69$, $SE = 0.23$, $z = 3.048$, $p = 0.0023$). There were no significant differences between New_Go_Prefix items (mean = 0.65 ± 0.062) and New_Me_Prefix items (mean = 0.63 ± 0.061) ($\beta = 0.069$, $SE = 0.22$, $z = 0.31$, $p = 0.76$), indicating that participants were not more likely to generalize to the novel direction (new prefix) for non-alternating [go] items than alternating [me/mo] items.

Because participants successfully generalized the alternating affix from suffixes to

prefixes, but not vice-versa, a pairwise comparison between the *New_Me_Prefix* for the *SuffixOnly* condition and the *New_Me_Suffix* for the *PrefixOnly* condition was made. This comparison was not statistically significant ($\beta = -0.38$, $SE = 0.33$, $z = -1.18$, $p = 0.24$), suggesting that even though only participants in the *SuffixOnly* condition generalized alternating suffixes to prefixes, there were no overall differences in generalization of the alternating affixes. In addition, the interaction between *New_Me_Prefix* and *New_Me_Suffix* by *PrefixOnly* and *SuffixOnly* was not statistically significant ($\beta = -0.24$, $SE = 0.29$, $z = -0.82$, $p = 0.42$).

3.5 Discussion

The results of Experiment 1 suggested that participants were able to learn the behaviors of both alternating and non-alternating affixes when they were exposed to each type of affix with the same frequency. Participants were more likely to generalize to novel affixes for non-alternating items than alternating items. The patterns of alternating affixes may be harder to learn than patterns of non-alternating affixes (exceptions). However, learners did not show preferences for a direction of spreading.

4 Experiment 2

In Experiment 2, participants were trained on a vowel harmony pattern that is ambiguous between directionality (e.g., stem to suffix harmony) and exceptionality (e.g., [me/mo] alternates while [go] fails to alternate). If participants infer that the non-alternating affix indicates that harmony does not apply in that direction, participants will fail to generalize the alternating affix to the novel direction. However, if participants infer that the non-alternating affix is an exception, participants will generalize the alternating affix to the novel direction.

4.1 Participants

All thirty-six participants were adults who were fluent in English speakers recruited from the psychology departments research familiarization pool at a private institution in the Pacific Northwest of the USA. No participants had any previous experience with a vowel harmony system, natural or artificial, nor did they participate in Experiment 1. Participants were randomly assigned to one of two training conditions: *SuffixAlternateOnly* and *PrefixAlternateOnly*.

SuffixAlternateOnly	PrefixAlternateOnly
bono bonomo gobono	bono mobono bonogo
deke dekeme godeke	deke medeke dekego
nepe nepeme gonepe	nepe menepe nepego

Table 3: Examples of Training Items in Experiment 2.

4.2 Design and Materials

The stimuli of Experiment 2 were identical to Experiment 1 with the following differences. In the *SuffixAlternateOnly* condition, the alternating affix was the suffix, and the non-alternating affix was the prefix (e.g., [bide, bideme, gobide]); in the *PrefixAlternateOnly* condition, the alternating affix was the prefix, and the non-alternating affix was the suffix (e.g., [bide, mebide, bidego]). See Table 3 for examples of training stimuli. The test items in Experiment 2 were identical to those found in Experiment 1.

4.3 Procedures

The procedure of Experiment 2 was identical to the Experimental condition of Experiment 1.

4.4 Results

Proportion of correct responses (means and standard errors of the mean) for all test items are given in Figure 2. The data were analyzed in the same manner as Experiment 1.

Participants in both conditions were overall more likely to select the correct ([go]) forms for both prefixes (mean = 0.76 ± 0.042) and suffixes (mean = 0.80 ± 0.039), than the correct ([me/mo]) forms for alternating prefixes (mean = 0.63 ± 0.048) and suffixes (mean = 0.70 ± 0.032) ($\beta = 0.69$, $SE = 0.34$, $z = 2.02$, $p = 0.043$), based on a model with [me/mo] items set as the baseline, with random intercepts for Subjects and Items and random slopes for Subjects by test item. This model showed a significant intercept ($\beta = 0.93$, $SE = 0.23$, $z = 4.07$, $p < 0.001$), suggesting that participants learned the behavior of the alternating affix, even though there were significantly more correct responses to non-alternating affixes.




Figure 2. Experiment 2 results (means and confidence intervals of the mean).

The model with New_Me_Suffix and the PrefixAlternateOnly condition items as the baseline showed a significant intercept, suggesting that participants in the PrefixAlternateOnly condition generalized the alternating affixes from prefixes to suffixes, $\beta = 0.86$, $SE = 0.25$, $z = 3.47$, $p < 0.001$. There was no significant difference between correct responses to New_Go_Prefix (mean = 0.72 ± 0.066) items compared to New_Me_Suffix (mean = 0.68 ± 0.067) ($\beta = 0.25$, $SE = 0.28$, $z = 0.87$, $p = 0.39$), showing no bias to extend the non-alternating affixes to a novel direction compared to alternating affixes.

The model with New_Me_Prefix items and the SuffixAlternateOnly condition as the baseline showed a significant intercept, suggesting that participants in the SuffixAlternateOnly condition generalized the alternating affixes from suffixes to prefixes ($\beta = 0.48$, $SE = 0.24$, $z = 2.02$, $p = 0.043$). There was a significant difference between New_Go_Suffix items (mean = 0.74 ± 0.061) and New_Me_Prefix items (mean = 0.60 ± 0.069) ($\beta = 0.73$, $SE = 0.28$, $z = 2.64$, $p = 0.0083$), indicating that participants were more likely to generalize to the novel direction for non-alternating [go] items than alternating [me/mo] items.

A pairwise comparison between the New_Me_Prefix items for the SuffixAlternateOnly condition and the New_Me_Suffix items for the PrefixAlternateOnly condition showed no significant differences between the two conditions ($\beta = 0.39$, $SE = 0.34$, $z = 1.13$, $p = 0.26$). These results suggest that there were no overall differences in generalization of the alternating affixes to novel affixes. In

addition, the interaction between `New_Me_Prefix` and `New_Me_Suffix` by `PrefixAlternateOnly` and `SuffixAlternateOnly` was not statistically significant ($\beta = -0.41$, $SE = 0.32$, $z = -1.29$, $p = 0.20$).

4.5 Discussion

Participants in Experiment 2 showed successful generalization of the alternating affix to a novel affix for both prefixes and suffixes. When participants were shown a non-alternating affix in prefix position, and an alternating affix in suffix position, participants generalized the alternating affix to the prefix position (and vice versa for suffixes). This suggests that participants inferred that the non-alternating affix was an exception, rather than an indicator of the directionality of the harmony pattern. Because there were no significant differences in generalization of the alternating affix to prefixes or suffixes, there was no evidence of a directionality bias for vowel harmony. Instead, participants inferred a bi-directional harmony pattern with exceptional non-alternating affixes.

5 General Discussion

The present study showed the results from two artificial language learning experiments with a front/back vowel harmony pattern in which one affix (either prefix or suffix, depending on the condition) alternated between [me] and [mo] depending on the phonetic feature of the stem vowels. In both experiments, participants were more likely to select the correct ([go]) forms for both prefixes and suffixes, than the correct ([me/mo]) forms for alternating affixes.

In Experiment 1, participants were able to learn the patterns of both alternating and non-alternating prefixes and suffixes, but only participants in the `SuffixOnly` condition generalized the alternating affix to a novel direction; participants generalized to novel alternating prefixes but not alternating suffixes. This result could be interpreted as support for a right-to-left bias in vowel harmony but should be taken with caution for two reasons. First, there were no statistically significant differences between the `PrefixOnly` and `SuffixOnly` conditions. Second, overall learning of the alternating affix in the `PrefixOnly` condition was relatively weak, even for trained affixes, which may account for the failure to generalize to suffixes.

In Experiment 2, participants were exposed to training data that contained non-alternating affixes in prefix position while alternating affixes were all suffixes, or vice versa. Participants extended the alternating affix to both prefixes and suffixes, suggesting that participants inferred that the non-alternating affix was an exception to the phonological rule rather than an indicator of directionality. Moreover, having non-alternating affixes in prefix position in the training phase did not reduce the amount of generalization to the novel direction. In other words, learners assumed that the [go] affix should also be non-alternating in a new location.

Overall, it appears that the patterns of alternating affixes may be harder to learn than patterns of non-alternating affixes (i.e., exceptions). This is not surprising, as alternating affixes involve more variation than non-alternation. Additionally, learners did not show preferences for a direction of generalization when presented with items spreading in opposite directions. These results suggest a learning bias towards non-alternation over alternation, which are in line with previous findings (e.g., Coetzee 2009, White 2014, Baer-Henney et al. 2015). Another possibility is that learning a vowel harmony language with both alternating and non-alternating affixes may be easier when the alternating and non-alternating affix share different roles (e.g., prefixes and suffixes respectively), as correct responses in Experiment 2 were numerically higher than Experiment 1.

While participants showed more correct responses to non-alternating affixes compared to the alternating affixes, it is possible that participants could have shown a bias towards vowel harmony even for non-alternating affixes. Because half of the test items had a back vowel stem, half of the items were consistent with both non-alternation and vowel harmony. Finley (to appear) showed that learners are biased towards harmony

in non-alternating affixes. Exploratory analyses suggests that this is the case in the present study; overall means for *New_Go_Suffix* items was 0.87 for back vowel (harmonic) stems, and 0.66 for front vowel (disharmonic) stems, and overall means for *New_Go_Prefix* items was 0.80 for back vowel (harmonic) stems, and 0.68 for front vowel (disharmonic) stems.

Our results also suggested that there were no learning biases towards directionality in vowel harmony, suggesting that learners infer a bi-directional harmony pattern when exposed to learning data that are ambiguous between a bi-directional pattern with exceptions, and a single-direction (e.g., prefix-to-stem or stem-to-prefix) pattern. These results help to better understand how learners make use of ambiguous information in learning novel linguistic patterns.

It is also important to keep in mind the limitations of the current study. The experiments involved a root-controlled harmony pattern, which means that the stem vowels determined the phonetic features of the affix vowels. In dominant-recessive and affix-controlled harmony, which our study did not test, the typological asymmetry between prefixes and suffixes is more evident (Bakovic 2000). It remains a possibility that learners might have different responses to novel items if the training data were designed to be affix-controlled or dominant-recessive. Moreover, it is unknown whether learners view the training data as stem-controlled, as the prefix could mirror the vowel structure of the stem. The fact that we failed to observe a directionality effect might be owing to those potential factors.

6 Conclusion

The present study tested how English-speaking adults learn novel phonological patterns with both alternating and non-alternating affixes in various affix roles (i.e., prefix and suffixes). Participants were able to learn the behavior of the alternating and non-alternating affixes. However, participants were more likely to generalize to novel affixes for non-alternating items. Learners assumed that the affix [go] should be non-alternating in a new location. This suggests a bias for non-alternation over alternation, which are consistent with past findings (e.g., Coetzee 2009, White 2014, Baer-Henney et al. 2015). Additionally, our results suggested that there were no learning biases towards directionality in vowel harmony (White et al. 2018), indicating that speakers infer a harmony pattern as bi-directional based on the prefixes or suffixes, consistent with the findings of Finley and Badecker (2009). These results help to better understand how learners make use of and encode ambiguous information in learning novel linguistic patterns.

While our findings did not suggest a learning bias towards directionality, it is possible that learners might respond to novel items differently if they are exposed to affix-controlled training data instead of stem-controlled training data. Future research will work to develop different training paradigms to understand how learners infer directionality in vowel harmony languages.

References

- Baković, Eric. 2000. Harmony, dominance and control. Doctoral dissertation, Rutgers University.
- Baer-Henney, Dinah, Frank Kügler, and Ruben van de Vijver. 2015. The interaction of language-specific and universal factors during the acquisition of morphophonemic alternations with exceptions. *Cognitive Science* 39: 1537–1569.
- Bates, Douglas, Martin Mächler, Ben Bolker, and Steve Walker. 2015. Fitting linear mixed-effects models using `{lme4}`. *Journal of Statistical Software* 67: 1–48.
- Coetzee, Andries. 2009. Learning lexical indexation. *Phonology* 26: 109-145.
- Coetzee, Andries, and Joe Pater. 2011. The place of variation in phonological theory. In *The Handbook of Phonological Theory* (2nd ed.), eds. J. Goldsmith, J. Riggle, and A. Yu, 1-38, Blackwell Publishers.

- Cutler, Anne, John Hawkins, and Gary Gilligan. 1985. The suffixing preference: A processing explanation. *Linguistics* 23: 723-758.
- Culbertson, Jennifer, Paul Smolensky, and Geraldine Legendre. 2012. Learning biases predict a word order universal. *Cognition* 122: 306-329.
- FindingFive Team (2019). FindingFive: A web platform for creating, running, and managing your studies in one place. FindingFive Corporation (nonprofit), NJ, USA. <https://www.findingfive.com>
- Finley, Sara. 2010. Exceptions in vowel harmony are local. *Lingua* 120: 1549-1566.
- Finley, Sara. 2012. Typological asymmetries in round vowel harmony: Support from artificial grammar learning. *Language and Cognitive Processes* 27: 1550-1562.
- Finley, Sara. To appear. Learning exceptions in phonological alternations. *Language and Speech*.
- Finley, Sara, and William Badecker. 2008. Analytic biases for vowel harmony languages. In *Proceedings of the 27th West Coast Conference on Formal Linguistics*, eds. N. Abner, and J. Bishop, 168-176.
- Finley, Sara, and William Badecker. 2009. Right-to-left biases for vowel harmony: Evidence from artificial grammar. In *Proceedings of the 38th North East Linguistics Society Annual Meeting*, ed. A. Shardi, M. Walkow, and M. Abdurrahman, 269-282.
- Finley, Sara, and William Badecker. 2010. Linguistic and non-linguistic influences on learning biases for vowel harmony. In *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society*, eds. S. Ohlsson, and R. Catrambone, 706-711.
- Hayes, Bruce, and Zsusa Londe. 2006. Stochastic phonological knowledge: The case of Hungarian vowel harmony. *Phonology* 23: 59-104.
- Hulst, Harry van der. 2016. Vowel harmony. In *Oxford Research Encyclopedia of Linguistics (OnLine)*, ed. A. Mark, online.
- Hyman, Larry. 2002. Is there a right-to-left bias in vowel harmony? Paper presented at the 9th International Phonology Meeting, Vienna.
- Lenth, Russell. 2016. Least-squares means: The R package lsmeans. *Journal of Statistical Software* 69: 1-33.
- Lenth, Russell. 2018. *emmeans: Estimated marginal means, aka least-squares means*. [R package]. Version 1.5.2-1. <https://cran.r-project.org/web/packages/emmeans/>
- White, James. 2014. Evidence for a learning bias against saltatory phonological alternations. *Cognition* 130: 96-115
- White, James, René Kager, Tal Linzen, Giorgos Markopoulos, Alexander Martin, Andrew Nevins, Sharon Peperkamp, Krisztina Polgárdi, Nina Topintzi, and Ruben van de Vijver. 2018. Preference for locality is affected by the prefix/suffix asymmetry: Evidence from artificial language learning. In *NELS 48: Proceedings of the Forty-Eighth Annual Meeting of the North East Linguistic Society*, eds. S. Hucklebridge, and M. Nelson, 207-220.

Department of Psychology
 Pacific Lutheran University
 Tacoma, WA 98447
 anqiwa@plu.edu
 finleysr@plu.edu