

Learning Exceptions in Phonological Alternations

Sara Finley (finleysr@plu.edu)

Department of Psychology, Pacific Lutheran University
Tacoma WA, 98447, USA

Abstract

The present study explores learning phonological alternations that contain exceptions. Participants were exposed to a back/round vowel harmony pattern in which a regular suffix followed harmony, varying between /e/ and /o/ depending on the back/round phonetic features of the stem, and an exceptional suffix that was always /o/ regardless of the features of the stem vowel. Participants in Experiment 1 learned the behavior of both suffixes, but performance for the non-alternating suffix was higher when the suffix happened to adhere to vowel harmony. In Experiment 2, participants were exposed only to the same suffixes as Experiment 1, but the non-alternating suffix only appeared in harmonic contexts, creating ambiguity between exceptionality and alternation. Participants only correctly selected the non-alternating suffix when it appeared in a harmonic context. This suggests that learners are biased towards alternating harmony patterns, but require concrete evidence of non-alternation to learn the non-alternating suffix.

Keywords: statistical learning; vowel harmony; learning biases; exceptions.

Introduction

One of the major challenges for characterizing the formal properties of language is that the vast majority of patterns are rife with exceptions. Exceptions not only pose a challenge for generative linguistics, but potentially to a language learner, who must distinguish between accidental regularities and true patterns that contain exceptionality. This paper focuses on phonological alternations, which govern the patterns of sounds that make up individual words. A phonological alternation occurs when an underlying condition is met within the sounds of a word or phrase that causes the sounds to change accordingly.

The vast majority of phonological alternations show some kind of exceptionality (Coetzee & Pater, 2011). It is these cases of exceptions that pose a challenge to the learner, as the learner must sort out which instances follow the alternation and which do not. Because the learner starts off without any specific knowledge of the language, the learner must make use of a combination of the data and any prior biases for how languages must be structured in order to determine how patterns apply, despite the many exceptions.

This paper focuses on vowel harmony, phonological pattern in which adjacent vowels (ignoring consonants) must share the same value of a phonological feature (Clements, 1976). In vowel harmony, all vowels of a word agree with each other in terms of a specific phonetically based feature value, such as front/back or high/low. For

example, in a back/round vowel harmony system in which features are shared from the leftmost vowel rightwards, words that begin with a back/round vowel (e.g., [o], [u]) must only contain back/round vowels, and words that begin with a vowel that is front/unround (e.g., [i], [e]) must only contain front/unround vowels. This creates alternations that apply when a suffix is added to the stem (creating morphophonological alternations). For example, in a simplified analysis of Hungarian, the dative suffix has two versions (allomorphs) depending on the back features of the stem vowel. When the stem vowel is front, the suffix surfaces as /nɛk/, as in [tsi:m-nɛk] ‘address-DAT’, and when the stem vowel is back, the suffix surfaces as /-nɔk/, as in [ɔblɔk-nɔk] ‘window-DAT’ (Hayes & Londe, 2006; Ringen, 1988; Vago, 1976).

Hungarian vowel harmony is complicated by its many types of exceptions. Front unround vowels are considered neutral, ‘transparent’ vowels because they fail to trigger vowel harmony if there is a back vowel preceding the front vowel (e.g., [radi:r-nɔk] ‘eraser-DAT’). In this case, there is a principled way of knowing whether vowel harmony should apply (the features of the front vowel). Despite the principled nature of transparent vowels, there are several cases in which the presence of a neutral transparent vowel in a stem cannot be used to predict the feature of the suffix. For example, while most stems containing neutral vowels take a front vowel suffix (e.g., [tsi:m-nɛk] ‘address-DAT’), there are some stems that take a back vowel suffix (e.g., [hi:d-nɔk], ‘bridge-DAT’). Determining whether a stem will take a front vowel suffix or a back vowel suffix is largely determined by statistical regularities of the additional vowel features in the stem (such as height) (Hayes & Londe, 2006). While these statistical regularities are relatively accurate at predicting the vowel quality of the suffix, they are general tendencies, rather than exceptionless alternations. In addition, some stems appear to take both front and back vowel suffixes, meaning that applying the harmony pattern is optional in these cases. In optionality, a phonological alternation may or may not apply. In these cases, speakers will produce both a front vowel suffix and a back vowel suffix for the same stem (e.g., both [ɔrze:n-nɔk]/[ɔrze:n-nɛk] were found in Hayes and Londe’s 2006 web search). Finally, another form of exceptionality occurs when some suffixes fail to alternate to vowel harmony, despite having all of the phonological properties necessary to alternate. While the dative suffix in Hungarian (discussed above) alternates in to adhere to vowel harmony,

there are some suffixes in Hungarian that fail to alternate, such as the temporal suffix */-kor/* (e.g., [øtkor] ‘at five (o’clock)’) (Kenesei, Vago, & Fenyvesi, 2002). In this type of exceptionality, a form will fail to apply despite meeting the conditions for alternation. In this case, there is only one form of the suffix, rather than a suffix that alternates.

Because most phonological alternations show some aspect of exceptionality such as those found in Hungarian, understanding how exceptionality affects the learning process is extremely important. While the bulk of research on learning exceptionality has focused on morphological processes (Prasada & Pinker, 1993), there has been an increased interest in learnability of exceptionality in phonological patterns (Baer-Henney, Kügler, & van de Vijver, 2014; Coetzee, 2009). For example, Baer-Henney et al. (2014) found that learners were more likely to generalize a phonetically grounded vowel harmony pattern compared to a phonetically ungrounded disharmony pattern when the grounded pattern occurred in only 65% of the items. This suggests that learners have a bias towards harmony over disharmony, even when the input is maximally ambiguous between harmony and disharmony.

Unlike Baer-Henney et al. (2014), the present paper focuses on exceptions in phonological alternations in which a particular suffix fails to alternate in accordance with the phonological pattern in question. There are several reasons to focus on learning a language with both alternating and non-alternating morphemes. First, the morphologically based exceptions to phonological alternations are fairly common among languages, and can be replicated within an artificial grammar learning paradigm with relative ease. Second, several formal linguistic studies have explored the topic of non-alternating exceptions, providing important theoretical and bases for an experimental learning study (Zonfeld, 1978), specifically in vowel harmony (Finley, 2010).

Previous research on learnability of exceptional non-alternating forms has revealed a bias towards non-alternation over alternation, even when alternations occurred in 75% of the training data (Coetzee, 2009; Stave, Smolek, & Kapatsinski, 2013; White, 2014). However, it is possible that the bias towards non-alternation in Coetzee’s (2009) paper stems from a bias against alternations in stems (Beckman, 1998), rather a general non-alternation bias, as stimuli in this Coetzee’s study only involved stem alternations. Baer-Henney et al.’s (2014) study suggested a bias towards vowel harmony, which is largely driven by alternations. However, because Baer-Henney et al.’s (2014) were focused on naturalness, it is unclear if naturalness was the drove the bias towards harmony. The present study will help tease apart how learners respond to non-alternating affixes.

Previous research has shown that adult participants can learn a back/round vowel harmony pattern with relatively minimal exposure to the pattern (e.g., as pairs of words, [bodo-bodomu], [bede-bedemi], etc.) (Finley & Badecker, 2009). The present experiment extends these findings by

including both an alternating and a non-alternating suffix. In Experiment 1, learners were exposed to a stem+suffix pattern where one suffix alternated according to vowel harmony, while a second stem failed to alternate. In Experiment 2, participants were exposed to the same alternating and non-alternating suffixes, but were only shown the non-alternating suffix in a context consistent with vowel harmony. If participants show a bias towards non-alternation, participants should be able to learn the behavior of the non-alternating suffixes in both experiments, but if participants are biased towards harmony over disharmony, then participants should fail to learn the non-alternating suffix without direct exposure.

Experiment 1

In the present study, learners were trained on a novel vowel harmony pattern in which one suffix alternated between [-me] and [-mo] depending on the back/round features of the stem, and another suffix was [-go] regardless of the features of the stem.

Participants

All 26 participants were adults who were fluent English speakers recruited from the psychology subject pool at Pacific Lutheran University, a small liberal arts college in Western Washington, USA. No participant had any previous experience with a vowel harmony system, natural or artificial.

Design

Participants were exposed to a back/round vowel harmony pattern that was presented in a set of three words: stem, stem+suffix1 and stem+suffix2 (e.g., [kine, kinego, kineme]). For half of the participants, suffix1 was the alternating [-me/-mo] suffix, while the other half of participants heard the non-alternating [-go] suffix as suffix1. In the alternating ([-me/-mo]) suffix items, stems triggered a suffix vowel that was either /e/ or /o/ depending on whether the vowels in the stems contained front vowels (/i/ or /e/) or back vowels (/o/ or /u/). In the case of the non-alternating suffix ([-go]), the suffix vowel was always /o/ regardless of whether the stem vowels were front or back. All stems were of the form CVCV (e.g., [keti]) with the vowels following back/round harmony constraints (all stem vowels were either both front or both back, and never disharmonic), and the consonants drawn from the set (/p, t, k, b, d, g, m, n/). There were 24 triads: stem followed by stem+suffix1 followed by stem+suffix2. These 24 sets of items were presented eight times, each in a random order. Examples of the exposure stimuli can be found in Table 1.

Following exposure, participants were presented with a two-alternative, forced-choice test in which participants were asked to decide which was more likely to come from the language they had just heard. Participants heard two words, each identical except for the final vowel ending: either /e/ or /o/. Because the first two vowels obeyed harmony, the choice of the final vowel (/e/ or /o/) depended

on the back/round features of the vowels in the first two syllables. The test items were presented in a random order, with five test conditions, and 10 items for each condition. Old items were items that were heard in the training set, and contained a mixture of both [-go] and [-me/-mo] suffixes. New-go, and New-me/mo items contained stems not heard in the training set, but with the [-go] and [-me/-mo] suffixes, respectively.

Table 1: Examples of Exposure Stimuli.

Stem Vowel	Stem	Stem+Suffix [-me/-mo]	Stem+Suffix [-go]
Back	tunu	tunumo	tunugo
	muto	mutomo	mutogo
	podo	podomo	podogo
	tonu	tonumo	tonugo
Front	kine	kineme	kinego
	tepe	tepeme	tepego
	bimi	bimime	bimigo
	pedi	pedime	pedimo

Finally, there were Agglutinative test items in which the alternating ([-me/-mo]) suffix was affixed to a word containing a /stem-go/ word (e.g., [bidi-go-me] vs. [bidi-go-mo]). These items were designed to test whether learners based the harmony pattern on the vowel features of stem or on the vowel features of the preceding suffix. For example, if the learner chooses [bidi-go-me] over [bidi-go-mo], the learner has based harmony on the stem rather than the preceding suffix. On the other hand, if the learner chooses [bidi-go-mo] over [bidi-go-me], then the learner is basing harmony on the closest vowel to the suffix. Cross-linguistically, vowel harmony languages with exceptions tend to follow locality principles and choose the closest vowel as a source for harmony, rather than the stem, if the stem is disharmonic with the closest vowel (Finley, 2010). For example, in Turkish, the suffix [-gen] does not alternate with back/round vowel harmony, but suffixes that follow the [-gen] suffix agree with the non-alternating suffix, rather than the stem (e.g., [tʃok-gen-ler], *[tʃok-gen-lar] ‘polygons’). In Finley’s (2010) survey of non-alternating affixes in vowel harmony, no languages followed the non-local pattern. Thus, it is predicted that learners will be likely to follow the cross-linguistic locality principles in the Agglutinative test items. Examples of these test items can be found in Table 2; note that the correct item was always [-go] regardless of the harmonic context.

All stimuli were recorded by a male speaker of American English with some phonetics training in a sound attenuated booth at 12,000 Hz. Stress was placed on the first syllable using English pronunciation, with the exception that no vowels were reduced, meaning that some English syllables contained partial stress (as English reduces unstressed syllables). All stimuli items were normalized for intensity (set at 70dB), though participants were allowed to adjust

headphones to a comfortable volume during the experiment. All stimuli creation and modifications to sound files was performed in Praat (Boersma & Weenink, 2015). All phases of the experiment were run in Psyscope X (Cohen, MacWhinney, Flatt, & Provost, 1993). Participants were given both written and verbal instructions. The entire experiment took approximately 20 minutes to complete.

Table 2: Examples of Experiment 1 Test Stimuli;
* Indicates ungrammatical/nonlocal response

Items	Front Vowel Suffix vs. Back Vowel Suffix
Old	*bemege vs. bemego bemime vs. *bimimo
New-Go	*buduge vs. budugo *kipege vs kipego
New-Me/Mo	*budume vs. budumo kipeme vs. kipemo
Old-Agglut	*bimigome vs. bimigomo *bemegome vs. bemegomo
New-Agglut	*kipegome vs. kipegomo *midigome vs. midigomo

Results

Proportion of correct responses for all test items are given in Figure 1. We compared each test item to 50% chance via three separate one-sample *t*-tests. If participants learned the harmony pattern with the correct suffix alternations ([-me/-mo] following harmony and [-go] as non-alternating), learners should select the correct response significantly greater than chance. Participants were successful at selecting the correct response for Old items (mean = 0.77, SD = 0.20), $t(25) = 6.92, p < 0.001$, New-go items (mean = 0.76, SD = 0.21), $t(25) = 6.51, p < 0.001$, and New-me-mo items (mean = 0.68, SD = 0.24), $t(25) = 3.88, p = 0.001$, but not to Old Agglutinative items (mean = 0.43, SD = 0.26), $t(25) = -1.42, p = 0.17$ or New Agglutinative Items (mean = 0.42, SD = 0.29), $t(25) = -1.48, p = 0.15$. This suggests that learners were able to correctly identify when to follow the harmony pattern (with the [-me/-mo] suffix alternation) and when not to follow the harmony pattern (with the [-go] suffix), but that they did not show a bias towards the items that contained both the alternating and non-alternating suffixes.

It is possible that even though learners successfully selected the [-go] suffix at a rate higher than chance for New-go items, that this success was based on the stems containing a back vowel, which would obey harmony. If this were the case, then participants may not have learned the exceptionality of [-go], but simply applied regular vowel harmony to this suffix. To test this, we separated the responses to New-go items for front vowel stems and back

vowel stems. While both back vowel stem (harmonic) items (mean = 0.84, SD = 0.22), $t(25) = 7.87$, $p < 0.001$ and front vowel stem (disharmonic) items (mean = 0.69, SD = 0.32), $t(25) = 3.10$, $p = 0.005$ were significantly greater than chance, participants chose the [-go] suffix items significantly more when the stem vowels were back $t(25) = 2.12$, $p = 0.044$, indicating that participants learned non-alternation, but still showed a bias towards harmony. Note that this pattern of results did not hold for the alternating suffix: there was no significant difference between front and back stem vowel responses $t(25) = 0.56$, $p = 0.58$, and both back vowel stem items (mean = 0.70, SD = 0.28), $t(25) = 3.74$, $p = 0.001$ and front vowel stem items (mean = 0.66, SD = 0.32), $t(25) = 2.64$, $p = 0.014$ were significantly greater than chance.




Figure 1: Experiment 1 Results (Means and Standard Errors).

Discussion

The results of Experiment 1 suggest that learners are biased towards harmony over disharmony, but learners are able to learn the behavior of the non-alternating suffix. Experiment 2 extends these findings by exposing participants to a language in which the non-alternating suffix [-go] only appears in a harmonic context, creating ambiguity between alternation and non-alternation.

Experiment 2

In Experiment 2, participants were trained on a vowel harmony pattern with two suffixes, one which clearly followed a harmony pattern (alternating between [-me] and [-mo]) and the other ambiguous between a non-alternating morpheme (e.g., [-go] only appearing in a back vowel context).

Participants

All 13 participants were adults who were fluent English speakers recruited from the psychology subject pool at Pacific Lutheran University, a small liberal arts college in Western Washington, USA. No participant had any previous

experience with a vowel harmony system, natural or artificial, nor did they participate in Experiment 1.

Design

The design of Experiment 2 was identical to Experiment 1 with the following changes. The non-alternating [-go] suffix only appeared in harmonic (back vowel stem) contexts, which was achieved by eliminating all front vowel stem-go suffixed items from the stimuli. Because front vowel stem-go items were removed, it was no longer possible to present the items as a triad. Instead, all items were presented in pairs with the stem followed by the stem+suffix.

There were four different types of test items in Experiment 2. New-go, New-me/mo and New Agglutinative items were all identical to those found in Experiment 1. Old-go items tested whether learners inferred that the front vowel stems heard in training would take the disharmonic, non-alternating affix [-go]. These included several of the same items in Experiment 1 (e.g., *bemege vs. beme-go).

Results

Proportion of correct responses for all test items are given in Figure 2. We compared each test item to 50% chance via separate Bonferroni corrected one-sample t -tests.

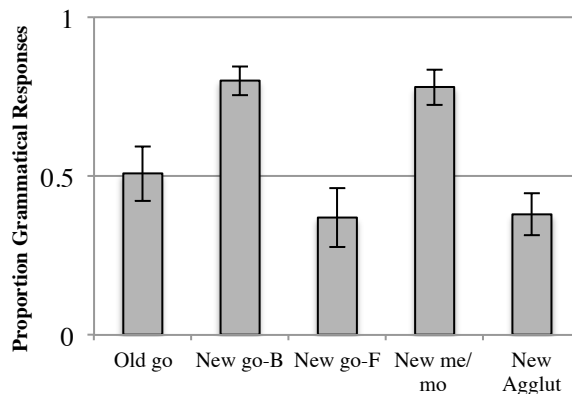


Figure 2: Experiment 2 Results (Means and Standard Errors).

If participants learned the harmony pattern with the correct suffix alternations ([-me/-mo] following harmony and [-go] as non-alternating), learners should select the correct response significantly greater than chance. Participants were not successful at inferring the non-alternating status of [-go] in front vowel stem items, as Old-go items (mean = 0.51 SD = 0.31), $t(12) = 0.089$, $p = 0.93$, were not significantly greater than chance. New-go items (mean = 0.58, SD = 0.21), $t(12) = 1.41$, $p = 0.175$, were also not significantly different from chance, suggesting that learners failed to learn the non-alternating status of the [-go] suffix in the ambiguous context.

New-me-mo items (mean = 0.78, SD = 0.20), $t(12) = 5.04$, $p < 0.001$ were significantly greater than chance, suggesting that participants were able to learn the status of

the alternating [-me/-mo] suffix. New Agglutinative items were also not significantly different from chance (mean = 0.38, SD = 0.24), $t(12) = -1.75$, $p = 0.10$, suggesting that learners may have a small bias towards non-local, stem controlled harmony in the presence of the disharmonic affix.

It is possible that while learners failed to successfully select the [-go] suffix at a rate greater than chance for New items, that this lack of success was based on the stems containing a front vowel, which would create a disharmonic sequence. To test this, responses to New-go items for front vowel stems were separated from back vowel stems. While back vowel stem (harmonic) items (mean = 0.80, SD = 0.16), $t(12) = 6.63$, $p < 0.001$ were significantly greater than chance front vowel stem (disharmonic) items (mean = 0.37, SD = 0.34), $t(12) = -1.41$, $p = 0.18$ were not, and trended towards harmony (as participants were lower than chance). Participants chose the [-go] suffix items significantly more when the stem vowels were back $t(12) = 4.94$, $p < 0.001$, indicating a bias towards harmony.

Because there were significantly different responses towards front vowel stems compared to back vowel stems for the [-go] suffix, it was important to compare the alternating [-me/-mo] suffix as well. Both back vowel stem items (mean = 0.75, SD = 0.30), $t(12) = 3.09$, $p = 0.009$ and front vowel stem items (mean = 0.82, SD = 0.24), $t(12) = 4.79$, $p < 0.001$ were significantly greater than chance, unlike the non-alternating [-go] suffix, there was no significant difference between front and back stem vowel responses $t(12) = 0.63$, $p = 0.54$, indicating a bias towards harmony that was not affected by exposure to the ambiguously non-alternating suffix.

Discussion

The present study showed the results from two artificial grammar learning experiments with a front/back vowel harmony pattern in which one suffix alternated between [-me] and [-mo] depending on the vowel quality of the stem vowels. In Experiment 1, participants learned the behavior of both the alternating and non-alternating suffixes, but were more likely to select the non-alternating suffix when the stem+suffix combination followed vowel harmony. In Experiment 2, participants were only exposed to forms containing the non-alternating suffix that obeyed harmony. At test, when presented with disharmonic instances of the stem-go items (with front vowel stems), participants were not significantly different from chance, but trended towards harmony for novel items. This suggests that learners inferred that the ambiguously non-alternating suffix was harmonic for back vowel stems, and optionally harmonic for front vowel stems. This bias towards patterns of harmony and optionality creates an overall pattern that is statistically more harmonic than disharmonic.

These results show an alternation bias, while previous results have shown a non-alternation bias (Coetzee, 2009; Stave et al., 2013; White, 2014). In the present study, learners were exposed to suffix alternations, while in previous work, learners were exposed to alternations within

a stem. Phonetic, phonological and psychological constraints work to prevent changes in the stem, while phonological processes work to promote changes (including underspecification of features) in suffixes. Pilot data has shown that the alternation bias found in the present experiments disappears when the stimuli are presented as single lexical items, rather than stem+suffix pairs, suggesting that the alternation bias is based in morphophonological alternations of affixes, which may be more likely to be underspecified.

Implications for Language Learning and Change

The results of the present study have implications for how languages with exceptional morphemes might be learned. If learners are biased towards optionality (and therefore a statistical harmony pattern) in ambiguous cases of exceptions, it is possible that learners will, over time, change a non-alternating morpheme to an alternating morpheme. Previous research on the evolution of vowel harmony languages in the Turkic language family cite several key motivators that favor and disfavor harmony (Harrison, Dras, & Kapcioglu, 2002), including frequency in the lexicon and specification in the lexicon. While a non-alternating morpheme may in principle seem simple (encoding [-go] rather than both [-ge] and [-go]), learners may prefer an underspecified lexicon in which the backness feature of the vowel is not stored, thus creating a simpler stored lexical representation. There is some evidence for overregularization of the non-alternating suffix [-ne] in Hungarian to [-na] following a back vowel suffix. This overregularization is particularly striking because it is said to last well into school years, and errors within vowel harmony systems are extremely rare (Dasinger, 1997). This suggests that language learners may show biases for harmony over non-alternation, which may drive languages to include more alternating suffixes. More research is needed to understand how this process might work, as changes in vowel inventory and loanword borrowings may cause the number of alternating harmony morphemes to decrease, rather than increase, which occurred in the harmony system of Korean (Finley, 2010). In addition, it is important to understand which phonological alternations may drive a bias towards alternations, and which alternations might drive non-alternations, as there is evidence that stem-based alternations may show a bias against alternations.

Agglutinative Test Items

The agglutinative test items were designed to probe whether learners show a bias for locality when presented with a front vowel suffix, followed by the non-alternating, disharmonic suffix [-go] and the choice between [-me] and [-mo]. Choosing [-me] indicates a preference for stem-controlled harmony, in violation of locality (as a the vowel [-go] intervenes between the stem vowel and the final suffix vowel). Choosing [-mo] indicates a preference for locality over stem-control, as the source for harmony is the closest

vowel is the non-alternating suffix. Both experiments trended towards non-locality, a violation of the locality principles described in Finley (2010). There are a few possible explanations for the effect. One possibility is that learners did not form a stable representation of the vowel harmony pattern. This lack of stability created a tie for the alternation between [-me] and [-mo] to occur in any other position besides a stem vowel. Another possibility is that the bias towards stem-control may interfere with any locality bias, showing no real bias within the data. A final possibility is that the way in which the items were spliced may have inadvertently favored the non-local items. Future research will include a control condition to investigate any initial biases that may be found within the stimuli.

Conclusions

The present study tested how adult, English speaking learners analyze exceptional, non-alternating morphemes. Participants were able to learn the behavior of the alternating morpheme, despite the presence of a disharmonic, non-alternating morpheme. However, participants were more likely to select the correct form of the non-alternating morpheme when it appeared in a harmonic context. This suggests a bias for harmony over non-alternation, which may have important implications for the evolution of languages with vowel harmony patterns. Interestingly, items that tested for a locality bias for the exceptions did not show any significant bias, and a slight trend toward non-locality, opposite of the typological descriptions shown in Finley (2010). It is possible that a clearer locality bias might be revealed for learners who acquire the non-alternating suffix, or for novel suffixes that were not presented as stem, stem+suffix pairs. Future research will work to understand how learners interpret non-alternating exceptions to vowel harmony, and how these exceptions interact with locality principles.

Acknowledgments

The author would like to thank participants at Waldorf College for providing valuable pilot data, as well as Ariel Goldberg, Neil Bardhan, Patricia Reeder, Paul Smolensky, William Badecker, Ashley Brandt Stine Lindvig and Kim Myhrvold. All errors are my own.

References

Baer-Henney, D., Kügler, F., & van de Vijver, R. (2014). The interaction of language-specific and universal factors during the acquisition of morphophonemic alternations with exceptions. *Cognitive Science*, 1–33. doi:10.1111/cogs.12209

Beckman, J. (1998). *Positional faithfulness*. University of Massachusetts, Amherst.

Boersma, P., & Weenink, D. (2015). Praat: Doing phonetics by computer.

Clements, G. N. (1976). The autosegmental treatment of vowel harmony. In W. Dressler & O. E. Pfeiffer (Eds.), *Phonologica*. Innsbruck: Institut für Sprachwissenschaft.

Coetzee, A. W. (2009). Learning lexical indexation. *Phonology*, 26, 109–145.

Coetzee, A. W., & Pater, J. (2011). The place of variation in phonological theory. In J. Goldsmith, J. Riggle, & A. Yu (Eds.), *The handbook of phonological theory (2nd ed.)*. Blackwell Publishers.

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

Dasinger, L. (1997). Issues in the acquisition of Estonian, Finnish, and Hungarian: A crosslinguistic comparison. In D. I. Slobin (Ed.), *The crosslinguistic study of language acquisition, Vol. 4*. Mahwah, New Jersey: Lawrence Erlbaum Associates Publishers.

Finley, S. (2010). Exceptions in vowel harmony are local. *Lingua*, 120(6), 1549–1566.

Finley, S., & Badecker, W. (2009). Right-to-left biases for vowel harmony: Evidence from artificial grammar. In A. Shardl, M. Walkow, & M. Abdurrahman (Eds.), *Proceedings of the 38th North East Linguistic Society Annual Meeting (Vol. 1, pp. 269–282)*. Amherst, MA: GLSA.

Harrison, K. D., Dras, M., & Kapicioglu, B. (2002). Agent-Based Modeling of the Evolution of Vowel Harmony. In *Proceedings of NELS (Vol. 32)*.

Hayes, B., & Londe, Z. (2006). Stochastic phonological knowledge: The case of Hungarian vowel harmony. *Phonology*, 23, 59–104. doi:10.1017/S0952675706000765

Kenesei, I., Vago, R. M., & Fenyvesi, A. (2002). *Hungarian*. London: Routledge.

Prasada, S., & Pinker, S. (1993). Generalizations of regular and irregular morphology. *Language and Cognitive Processes*, 8, 1–56.

Ringen, C. O. (1988). Transparency in Hungarian vowel harmony. *Phonology*, 5(2), 327–342.

Stave, M., Smolek, A., & Kapatsinski, V. (2013). Inductive bias against stem changes as perseveration: Experimental evidence for an articulatory approach to output-output faithfulness. In *Proceedings of the 35th Annual Proceedings of the Cognitive Science Society* (pp. 3454–3459).

Vago, R. M. (1976). Theoretical implications of Hungarian vowel harmony. *Linguistic Inquiry*, 7, 243–263.

White, J. (2014). Evidence for a learning bias against saltatory phonological alternations. *Cognition*, 130(1), 96–115.

Zonnefeld, W. (1978). *A formal theory of exceptions in generative phonology*. Lisse: Peter de Ridder Press.