The Role of Anti-harmony in Learning Neutral Vowels*

Sara Finley

Pacific Lutheran University

1 Introduction

Learning human language requires the ability to make sense of sparse data with many alternative interpretations. One possible explanation for how learners achieve this feat is that learners make use of learning biases. These learning biases may be domain general in nature, helping learners to find statistical regularities in the input in order to form abstract rules (Blevins 2004; Christiansen and Chater 2008). Domain specific learning biases may help learners to form patterns that are phonetically and phonologically natural (Wilson 2006; Finley and Badecker 2008; Moreton 2008; Culbertson 2012; Yin and White 2018). Language learning biases may also reflect an interaction of domain general processes with domain specific effects (Culbertson and Kirby 2016). For example, a general bias for simplicity over complexity (Moreton and Pater 2012; Moreton et al. 2017) could bias learners towards phonological patterns that are simplified. This can help to explain the fact that more complex patterns are more likely to be cross-linguistically rare or impossible, and are often more difficult to learn (Heinz et al. 2011; Lai 2015; McMullin 2016; Heinz 2018). Understanding the nature of these biases and how they shape linguistic typology is thus a great source of interest in linguistics, particularly in phonology, where phonological patterns are phonetically grounded, as well as structurally abstract.

One particularly promising phonological pattern to serve as a window into the cognitive and linguistic mechanisms that underlie language learning and language typology is vowel harmony. Vowel harmony is a phonological process whereby vowels in a given domain are required to share a phonological feature. For example, in Hungarian back vowel harmony, the dative suffix alternates between [-nɛk] and [-nɔk] depending on the vowel features of the stem (Robert M Vago 1976; Ringen 1980; Ringen 1988; Ringen and Vago 1998) (e.g., [yʃt-nɛk] 'cauldron-DAT' vs. [ɔblɔk-nɔk] 'window-DAT' (Hayes and Londe 2006)). Because vowel harmony is a relatively common phenomenon found across a wide range of languages, it has received a great deal of attention in terms of its typology and representations. One of the great pieces of interest within the study of vowel harmony is the representation of neutral vowels, particularly transparent neutral vowels.

Most languages with vowel harmony have some number of vowels that fail to participate in harmony. These neutral, or non-participating vowels, typically come in two varieties: opaque and transparent vowels. Opaque vowels 'block' vowel harmony, and create a new harmonic domain. In Turkish, low vowels are opaque to rounding harmony, and low vowel stems trigger a suffix that agrees in rounding with the low vowel (e.g., [buz-a-mu], *[buz-a-mu] 'is it ice-DAT') (Gafos and Dye 2011). Transparent vowels, on the other hand, are 'invisible' to the harmonic domain, and agreement of features passes through the neutral vowel. For example, in the Hungarian word [pɔlle:r-nɔk] 'foreman-DAT' (Hayes and Londe 2006), the initial back vowel selects the back vowel suffix [-nɔk], even though the front vowel in the stem is adjacent to the suffix vowel; harmony skips the neutral vowel [e]. While both transparent and opaque vowels appear widely cross-linguistically, transparent vowels tend to complicate phonological analyses (Baković and Wilson 2000; Gafos and Dye 2011).

Transparent neutral vowels typically require more complex representations because transparent vowels violate the assumption of locality (e.g., even when the neutral vowel is the closest stem vowel to the suffix, the neutral vowel does not determine the vowel quality of the suffix) (Gafos and Dye 2011). While native speakers appear to be able to learn this violation of locality, adult learners in an artificial language learning

^{*} The author would like to thank participants at Pacific Lutheran University, as well as Shelan Porter, Nicole Wirth, Liza Radford, Stella Wang, Sara Charania, and Barbara Millward. All errors are my own.

task required more exposure to learn a transparent vowel compared to an opaque vowel (Finley 2015a). This suggests that the non-local representations required for transparent vowels may be more complex, and therefore take longer to learn, but are well within the scope of what is learnable. More work is needed to understand when and how learners are able to cope with the complex representations and locality violations that transparent vowels bring.

One possibility is that while the representation of transparent vowels may be more complex, there may be other properties of transparent vowels that help learners discover their transparent nature. For example, high vowels are more likely to be transparent to back vowel harmony than low vowels (Hayes and Londe 2006; Smith 2016). Rebrus and Törkenczy (2015) hypothesize that this height effect is a result of the phonological representation of transparent vowels in back vowel harmony. High vowels are phonetically more likely to be 'in-between' front and back, which allows them to pattern with both front and back vowels, thereby allowing for both harmony and disharmony. This may provide an explanation for why front transparent vowels tend to be produced as more 'back' when in back vowel contexts than front vowel contexts (Benus and Gafos 2007). In addition, the 'in-between' representations for transparent vowels allows for anti-harmony. Anti-harmony occurs when a transparent vowel selects for a disharmonic affix. For example, in Hungarian, there are stems that contain a front neutral vowel that always select for a back vowel suffix (e.g., [fi:r-ok] 'tomb-PL' (Rebrus and Törkenczy 2015)). Rebrus and Törkenczy (2015) predict that only languages with transparent vowels allow for anti-harmony; languages with opaque vowels will not show anti-harmony. This prediction appears to be borne out for the known typology of back vowel harmony languages. Rebrus and Törkenczy (2015) explain the lack of anti-harmony in opaque vowels as a result of the representation of opaque vowels. While transparent vowels are 'in-between' front and back, phonetically, opaque vowels are clearly either front, or back, and therefore can only select for suffixes that share the feature value of the opaque vowel.

If Rebrus and Törkenczy's (2015) predictions are correct, it may be possible that learners might infer that a language has transparent vowels, rather than opaque vowels, merely through exposure to antiharmony. The presence of anti-harmony, while seemingly making transparent vowels more complex than opaque vowels, could potentially inform the learner that the representation of the neutral vowel is 'in between' front and back (for back vowel harmony), and will behave transparently.

Rebrus and Törkenczy's (2015) predictions can be tested using an artificial language learning paradigm (Finley 2015b). In this experimental paradigm, learners are exposed to a vowel harmony pattern, and then tested on their learning of that pattern. Artificial language learning experiments have been relatively successful at teasing apart predictions about linguistic theory and learnability (Wilson 2006; Moreton 2008; Finley and Badecker 2009; Culbertson et al. 2012; Culbertson 2012; Finley and Badecker 2012; White 2014; Culbertson and Newport 2015; Fedzechkina et al. 2016; Finley 2017b; Finley 2017a). In addition, language learners have shown relative success at learning and generalizing vowel and consonant harmony patterns in an laboratory setting (Pycha et al. 2003; Wilson 2003; Finley and Badecker 2008; Moreton 2008; Finley 2011; Finley 2012; Finley and Badecker 2012; McMullin and Hansson 2014; Finley 2015c; Finley 2015a; Lai 2015; Kimper 2016; Finley 2017a; White et al. 2018). The results of these studies generally show that learners are sensitive to the phonetic, phonological, and structural representations that shape the typology of vowel and consonant harmony languages; learners are more likely to learn and generalize to a novel harmony pattern when it fits the typology of vowel harmony languages. This suggests that if Rebrus and Törkenczy's (2015) model is correct, its predictions may be detectable in an artificial language learning experiment.

In the present study, learners were exposed to a novel vowel harmony language with a neutral vowel that either behaved in accordance with anti-harmony or did not adhere to anti-harmony. In addition to regular vowel harmony (without the neural vowel), participants were trained on stems that contained only neutral vowels, and then tested on their inferences about whether the neutral vowel was transparent or opaque. This 'Poverty of the Stimulus' paradigm (Wilson 2006) allows the experimenter to test for learning biases through inferences, allowing the researcher to create a more streamlined artificial language that is easy to learn in a single training session (as previous research on neutral vowels in vowel harmony showed that learning transparent vowels is relatively difficult (Finley 2015a)). It is hypothesized that if learners are exposed to a neutral vowel that conforms to anti-harmony, then learners will be more likely to infer that the neutral vowel is transparent.

2 Method

2.1 *Participants* Sixty-eight adult American English-speaking participants at Pacific Lutheran University, a small liberal arts college in the Pacific Northwest of the United States, received course credit for their participation. No participant had prior exposure to a vowel harmony language (artificial or natural). There were 23 participants in the Anti-Harmony condition and the Neutral-e conditions, and 22 participants in the Neutral-o condition.

2.2 Design Participants were exposed to one of three vowel harmony languages, adapted from Finley (2015a). The language was a back/round vowel harmony language in which the back/round features of the CVCVC stems determined the value of a suffix (either [-o] or [-e]). When stem vowels were front ([e], or [i]), the suffix was always the front vowel [-e]. When stem vowels were back ([o], or [u]), the suffix vowel was always the back vowel [-o]. When the stem contained the neutral vowel [a], the suffix vowel varied depending on the training condition. In the Neutral-o Condition, the suffix vowel was always [-o]. In the Neutral-o Condition, the suffix vowel was always [-o]. In the Neutral-e Condition, the suffix vowel was always [-o] for half of the stems, and [-e] for the other half of stems.

The low vowel [a] was chosen as a neutral vowel for several reasons. First, previous research in vowel harmony learning with English speakers showed that learners naturally treat the low yowel as neutral in a back/round harmony language (Finley and Badecker 2009). Second, high vowels are more likely to be transparent in back vowel harmony languages (Hayes and Londe 2006; Smith 2016). This means that if learners are biased towards transparency, it cannot be explained simply by the vowel in question (though the height of the vowel did not seem to affect learning in Finley 2015a). Third, the low vowel [a] is phonetically centralized, meaning that this vowel may have a representation 'in-between' front and back vowels, which, according to Rebrus and Törkenczy (2015), should make [a] a good candidate to be a neutral vowel. One problem with using [a] as the neutral vowel is that the artificial harmony pattern is a 'color' harmony (back and round vowel harmony), and not a back vowel harmony language. Because the experiment makes use of English speakers and an English vowel inventory, it is impossible to induce a contrast between front and back vowels without conflating backness and rounding (as English does not contain front round vowels). Because Rebrus and Törkenczy's (2015) predictions are about back vowel harmony, and not round harmony, it is impossible to fully test predictions based on back harmony, as the exposure language was not a strict back vowel harmony language. Because the language is a back/round harmony language, it is unclear whether participants will interpret [-e] as the harmonic suffix, patterning with unround vowels, or [-o] as the harmonic suffix, patterning with back vowels. For this reason, conditions are labeled as Neutral-e and Neutral-o, and Anti-Harmony. It is important to note that if learners interpret the neutral vowel as unround, that the Neutral-o condition would also adhere to anti-harmony, as there are languages with transparent vowels in which the stems with neutral vowels always select for the disharmonic vowel, such as Uyghur (Vaux 2000) (and likewise for the Neutral-e condition if learners interpret the neutral vowel as back).

Evidence of learning was assessed using a two-alternative forced-choice test that pitted one item ending in the [-e] suffix against another item (with the same CVCVC stem) ending in the [-o] suffix (e.g., [bipene] vs. [bipeno]). There were four different types of test items: Old (which were identical to those in the training set), New-Harmonic (which contained stems with [i, e, o, u] only), New-Front-Neutral (which contained a front vowel followed by a neutral vowel), and New-Back-Neutral (which contained a front vowel followed by a neutral vowel). If participants have inferred that the neutral vowel is transparent, then participants will be more likely to select the back vowel suffix when the initial vowel is back, compared to when the initial vowel is front. If the presence of anti-harmony aids the inference of transparent vowels, then learners exposed to anti-harmony will be more likely to infer that the neutral vowel is transparent than participants who were not exposed to anti-harmony.

2.3 *Stimuli* All stimuli were naturally produced by a female native English speaker from the Pacific Northwest who was unaware of the hypothesis of the experiment. All sound editing was done using Praat software (Boersma and Weenink 2017). Audio files were normalized to 70dB and screened for perceptibility.

The training set contained 18 stem+suffix pairs (e.g., [pideg pidege]). These items were evenly

distributed between stem-vowel type: six stems containing only front vowels, six stems containing only back vowels, and six stems containing only the neutral vowel [a]. Examples of the suffixed items for the 18 training items can be found in Table 1 below. The training items were identical for all conditions, except for the items containing the neutral vowel.

Condition	Front Vowel Stem	Back Vowel Stem	Neutral Vowel Stem	
Neutral-o	pidege	butono	banamo	
	mebite	gopubo	dakago	
Neutral-e	pidege	butono	baname	
	mebite	gopubo	dakage	
Anti-Harmony	pidege	butono	baname	
	mebite	gopubo	dakago	
Table 1: Examples of Training Stimuli				

The test contained 10 of each of the four types of test items discussed above (Old, New-Harmonic, New-Front-Neutral, and New-Back-Neutral). The set of test items was identical for all participants in all conditions. Items were counterbalanced so that the items ending in [e] occurred first about as often as the items ending in [o], and the 'correct' item appeared first about half of the time. All items appeared in a different random order for each participant. Examples of the 40 test items can be found in Table 2, below.

Test Condition	Front Vowel Affix	Back Vowel Affix	
Old	baname	banamo	
	mebite	*mebito	
New-Harmonic	bipene	*bipeno	
	*kupuge	kupugo	
New-Front-Neutral	bitame	bitamo	
	mepane	mepano	
New-Back-Neutral	bopane	bopano	
	nubade	nubado	

 Table 2: Examples of Test Stimuli (* Indicates Ungrammatical for all Conditions)

2.4 *Procedure* The experiment was presented on PsyScope X software (Cohen et al. 1993) on a Macintosh computer. Instructions were presented in a written format, but participants were also given basic verbal instructions before beginning the study. Participants heard the 18 training stimuli in a random order, repeated 10 times, with a different random order on each iteration. The test was also presented in a random order, and immediately followed the training. The entire procedure took approximately 15-20 minutes to complete. A copy of the PsyScope files, along with relevant stimuli files can be found at: https://osf.io/sc5ut/.

3 Results

Means and standard deviations for all participants can be found in figures 1 and 2. The data were fit into mixed effects logistic regression models with the Laplace approximation using the lme4 package (Bates et al. 2015) in R (R Development Core Team 2018). There were random intercepts for Subjects and Items, and random slopes for Training Condition by Subjects (as participants were randomly assigned to the various training conditions, but all other items were not randomly assigned). This was the maximal model justified by the design of the experiment that would reach convergence (Barr et al. 2013). In order to simplify the models, and to avoid failure for models to converge, the data were split into Old and New-Harmonic test items in one set of models, and New-Front-Neutral and New-Back-Neutral test items in a separate set of models. Analysis code can be found at: https://osf.io/sc5ut/.

3.1 Old and New-Harmonic Test Items Means and standard errors for the proportion of correct

responses for Old and New-Harmonic items can be found in Figure 1, below. Because the presence of a neutral vowel should not affect the behavior of stems not containing the neutral vowel, it was predicted that there should be no significant differences between conditions, as the presence or absence of anti-harmony should not affect the learnability of the general harmony pattern.



Figure 1: Proportion of harmonic responses: Old and New-Harmonic Items

While participants in the Anti-Harmony condition showed numerically lower correct responses to Old items, there was no significant difference in Old item responses between the Anti-Harmony (mean=0.55, SD=0.16), and the Neutral-o conditions (mean=0.61, SD=0.22), (β =0.34, SE=0.28, z=1.25, p=0.21), or between the Anti-Harmony and the Neutral-e conditions (mean=0.62, SD =0.16) (β =0.39, SE=0.25, z=1.56, p=0.12). There were no significant differences in New-Harmonic item responses between the Anti-Harmony (mean=0.60, SD=0.17), and either the Neutral-o (mean=0.57, SD=0.21), (β =-0.11, SE=0.27, z=0.42, p=0.67), or the Neutral-e conditions (mean=0.60, SD=0.20), (β =0.024, SE=0.26, z=0.093, p=0.93). While participants in the Anti-Harmony condition may have struggled somewhat to learn the specific items in the training set, this did not affect the learnability of the general harmonic pattern, as responses to New-Harmonic items were significantly above chance (β =0.42, SE=0.20, z=2.14, p=0.032).

3.2 New-Front-Neutral and New-Back-Neutral Test Items Means and standard errors for the proportion of back vowel suffix responses for New-Front-Neutral and New-Back-Neutral test items can be found in Figure 2, below. Each of the three training conditions was tested for a bias towards [a] as a transparent neutral vowel. If [a] is treated as a transparent neutral vowel, there should be significantly more back vowel suffix responses when the stem contains a back vowel (New-Back-Neutral items) compared to when the stem contains a front vowel (New-Front-Neutral items).



Figure 2: Proportion of back vowel suffix responses: New-Front-Neutral and New-Back-Neutral Items

In the Anti-Harmony condition, there were significantly more back vowel suffix responses for New-Back-Neutral items (mean=0.54, SD=0.17) compared to New-Front-Neutral items (mean=0.41, SD=0.14) (β =0.53, SE=0.19, z=2.79, p=0.0052), suggesting a bias towards a transparent neutral vowel. In the Neutral-o condition, there were significantly more back vowel suffix responses for New-Back-Neutral items (mean=0.67, SD=0.26) compared to New-Front-Neutral items (mean=0.53, SD=0.28) (β =0.70, SE=0.21, z=3.36, p<0.001), suggesting that participants in the Neutral-o condition were biased towards [a] as a transparent neutral vowel. In the Neutral-e condition, there were no significant differences between back vowel suffix responses for New-Back-Neutral items (mean=0.45, SD=0.25) compared to New-Front-Neutral items (mean=0.45, SD=0.25) compared to New-Front-Neutral items (mean=0.425, SD=0.21) (β =0.21, SE=0.19, z=1.13, p=0.26), suggesting no bias for [a] as a transparent neutral vowel. These results suggest that exposure to either back vowel suffixes or antiharmony, but not front vowel suffixes, lead participants to interpret [a] as a transparent vowel.

If participants in the Anti-Harmony condition showed more of a bias to interpret [a] as a transparent neutral vowel, then there should be a significant interaction between the Anti-Harmony conditions and each of the Neutral-o and Neutral-e conditions for the New-Neutral items, such that the difference between New-Front-Neutral and New-Back-Neutral items should be significantly greater in the Anti-Harmony condition. However, there were no significant interactions comparing the Anti-Harmony condition to either the Neutral-o condition (β =0.18, SE=0.28, z=0.63, p=0.53), or the Neutral-e condition (β =0.31, SE=0.27, z=1.16, p=0.25). In addition, the interaction between the Neutral-o and the Neutral-e conditions was only marginally significant (β =0.49, SE=0.28, z=1.72, p=0.085). This suggests that while the presence of anti-harmony and [o] as the suffix biased the learners towards [a] as a transparent neutral vowel, no condition had a greater bias towards [a] as the transparent vowel than any other.

4 Discussion

Participants were exposed to a novel back/round vowel harmony language with a neutral vowel [a]. During training, participants were only shown the neutral vowel in isolation (e.g., without any other vowels). Participants who were trained that the neutral vowel could co-occur with both the front and the back vowel suffix, or trained that the neutral vowel triggered a back vowel suffix, were biased to treat the neutral vowel as transparent. While participants who were trained that the neutral that the neutral that the neutral fall triggered the front vowel suffix [e] showed no bias for [a] as a transparent vowel.

If learners inferred that the neutral vowel [a] was unround, and therefore, given round harmony, should be grouped with the [e] suffix, then the Neutral-e condition would be considered the Harmonic condition, and the Neutral-o condition should be considered an Anti-Harmony condition. This is interpretation is not unreasonable; while [a] is a back vowel, it is both unround, and centralized in American English (Hillenbrand et al. 1995). If participants assumed that [a] should be grouped with [e] as an unround vowel, then predictions from Rebrus and Törkenczy (2015) hold; the two anti-harmony conditions showed biases for [a] as a transparent vowel, and the harmony condition did not. The results of the present experiment suggest that it is enough to expose learners to only a few instances of disharmony, as the Anti-Harmony condition only provided three instances of the [o] suffix and three instances of the [e] suffix, and this was enough to bias the learner towards interpreting the neutral vowel as transparent.

While participants in the Neutral-e condition did not show a bias towards [a] as a transparent vowel, it is not clear that there was a bias towards [a] as an opaque vowel, either. While responses to [o] as the appropriate suffix were significantly below 50% for New-Front-Neutral items (β =-0.38, SE=0.17, z=-2.27, p=0.023), responses to New-Back-Neutral items, were not significantly below 50% (but trended in that direction) (β =-0.16, SE=0.16, z=-0.98, p=0.325)). It is likely that participants in the Neutral-e condition did not have enough information to infer whether the vowel was transparent or opaque, since there are languages with transparent vowels that do not show anti-harmony (Rebrus and Törkenczy 2015). This means that even if learners can use anti-harmony as a cue to infer that the language has transparent vowels, learners must still be able to learn the behavior of transparent vowels, even without anti-harmony. Finley (2015a) suggested that a large number of items containing the transparent vowel can help learners infer that a neutral vowel is transparent.

In Finley's (2015a) experiments, the neutral vowel was always a front, unrounded vowel, while the present experiment used the low vowel [a]. Future research could replicate the present experiment with a neutral vowel similar to the those used in Finley (2015a), to test whether the reverse effect (i.e., a bias for transparent in the Neutral-e condition rather than a bias in the Neutral-o condition) would be obtained.

5.1 *Responses to Old Items* One curiosity about the present experiment is that participants in the Anti-Harmony condition responded relatively poorly to the Old items. One possibility could be that the participants failed to learn the suffix that belonged to the items that contained the neutral vowel, as these varied. However, when the three Old items that contained the neutral vowel were removed, the mean correct response was only 0.53, suggesting that participants struggled to remember the specific suffix attached to specific stems. It may be that learning that the language tolerates disharmony, even if only presented in the context of neutral vowels, makes participants more willing to accept disharmony in lexical items (e.g., the Old items). Note, however, that participants generally did learn the harmony pattern.

5.2 Learning Versus Generalization The present study made use of the 'Poverty of the Stimulus' paradigm (Wilson 2006); participants were exposed to neutral vowels in stems that contained only neutral vowels. This was done to determine a learner's inferences about the behavior of the neutral vowel in a front or back vowel context at test. A question one might ask is whether exposing participants to anti-harmony in a less ambiguous context might reveal more about whether anti-harmony serves as an aid to learning transparent vowels. In a pilot study, participants were exposed to one of four conditions: Opaque, Transparent, Opaque-Anti-Harmony, and Transparent-Anti-Harmony. These conditions contained training items similar to the 18 used in the present experiment, but also included 12 additional items demonstrating the role of the neutral vowel in a back vowel context (six items) or a front vowel context (six items). In the conditions without anti-harmony, stems with only neutral vowels always triggered [o], while the neutral vowel in the conditions with anti-harmony triggered either [e] or [o]. In the opaque vowel conditions, the neutral vowel always triggered the [o] suffix, while in the transparent vowel conditions, the initial vowel determined the vowel quality of the suffix. In general, participants were able to learn the behavior of the opaque vowel, but not the behavior of the transparent vowel, similar to findings in Finley (2015a). However, the learning problem was much easier for participants in the opaque conditions because the suffix vowel was always [o], while it varied in the transparent conditions. In order to demonstrate learnability of transparent compared to opaque vowels, one would need to ensure that the complexity of the learning problem was equivalent for all conditions (Moreton and Pater 2012; Moreton et al. 2017).

5 Conclusion

The present study used an artificial grammar learning paradigm to explore the prediction that exposure to anti-harmony might help learners infer that the neutral vowel in a vowel harmony language is transparent. Results indicated that exposure to a back/round harmony language with the neutral vowel selecting either back vowel suffixes, or both front and back vowel suffixes, could induce a bias towards transparent vowels. Assuming that participants inferred that the centralized [a] paired with [o] harmonically, then the predictions that exposure to anti-harmony could induce a bias towards a transparent vowel interpretation were borne out. However, the bias towards a transparent vowel was not significantly different for the anti-harmony conditions compared to the harmony condition, suggesting that this effect should be replicated with other neutral vowels.

References

- Baković, Eric, and Colin Wilson. 2000. Transparency, strict locality, and targeted constraints. In Proceedings of the West Coast Conference on Formal Linguistics 19, 43–56.
- Barr, Dale J., Roger Levy, Christoph Scheepers, and Harry J. Tily. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language* 68: 255–278. doi:10.1016/j.jml.2012.11.001.
- Bates, Douglas, Martin Mächler, Ben Bolker, and Steve Walker. 2015. Fitting linear mixed-effects models using {Ime4}. Journal of Statistical Software 67: 1–48.
- Benus, Stefan, and Adamantios I. Gafos. 2007. Articulatory characteristics of Hungarian 'transparent' vowels. *Journal* of *Phonetics* 35: 271–300.
- Blevins, Juliette. 2004. Evolutionary phonology: The emergence of sound patterns. Cambridge: Cambridge University Press.
- Boersma, Paul, and David Weenink. 2017. Praat: Doing phonetics by computer.
- Christiansen, Morten H., and Nick Chater. 2008. Language as shaped by the brain. *Behavioral & Brain Sciences* 33: 489–558.
- Cohen, Jonathan D., Brian MacWhinney, Matthew Flatt, and Jefferson Provost. 1993. PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments and Computers* 25: 257–271.
- Culbertson, Jennifer. 2012. Typological universals as reflections of biased learning: Evidence from artificial language learning. *Language and Linguistics Compass* 6: 310–329. doi:10.1002/lnc3.338.
- Culbertson, Jennifer, and Simon Kirby. 2016. Simplicity and specificity in language: Domain general biases have domain specific effects. *Frontiers in Psychology* 6.
- Culbertson, Jennifer, and Elissa L Newport. 2015. Harmonic biases in child learners: In support of language universals. *Cognition* 139: 71–82. doi:10.1016/j.cognition.2015.02.007.
- Culbertson, Jennifer, Paul Smolensky, and Geraldine Legendre. 2012. Learning biases predict a word order universal. *Cognition* 122: 306–329. doi:10.1016/j.cognition.2011.10.017.
- Fedzechkina, Mariya, Elissa L Newport, and T. Florian Jaeger. 2016. The miniature artificial language learning paradigm as a complement to typological data. In *The usage-based study of language learning and multilingualism*, ed. Tyler Ortega, L., H. I. A., Park, and M. Uno, 211–232. Washington, DC: Georgetown University Press.
- Finley, Sara. 2011. The privileged status of locality in consonant harmony. *Journal of Memory and Language* 65: 74–83.
- Finley, Sara. 2012. Typological asymmetries in round vowel harmony: Support from artificial grammar learning. Language and Cognitive Processes 27: 1550–1562. doi:10.1080/01690965.2012.660168.
- Finley, Sara. 2015a. Learning nonadjacent dependencies in phonology: Transparent vowels in vowel harmony. *Language* 91: 48–72. doi:10.1353/lan.2015.0010.
- Finley, Sara. 2015b. Consequences of monotonicity: Representations and learnability. *Theoretical Linguistics* 41. doi:10.1515/tl-2015-0003.
- Finley, Sara. 2015c. Learning exceptions in phonological alternations. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, 37:698–703. Austin, TX: Cognitive Science Society.
- Finley, Sara. 2017b. Learning metathesis: Evidence for syllable structure constraints. *Journal of Memory and Language* 92. doi:10.1016/j.jml.2016.06.005.
- Finley, Sara. 2017a. Locality and harmony: Perspectives from artificial grammar learning. *Language and Linguistics Compass* 11. doi:10.1111/lnc3.12233.
- Finley, Sara, and William Badecker. 2008. Analytic biases for vowel harmony languages. In *West Coast Conference of Formal Linguistics*, ed. Natasha Abner and Jason Bishop, 27:168–176.
- Finley, Sara, and William Badecker. 2009. Artificial language learning and feature-based generalization. Journal of Memory and Language 61: 423–437.
- Finley, Sara, and William Badecker. 2012. Learning biases for vowel height harmony. *Journal of Cognitive Science* 13: 287–327.
- Gafos, Adamantios I., and Amanda Dye. 2011. Vowel harmony: Transparent and opaque vowels. In *The Blackwell companion to phonology*, ed. Marc VanOostendoorp, Colin Ewen, Elizabeth Hume, and Keren Rice, 2164–2189.

Wiley-Blackwell.

- Hayes, Bruce, and Zsuzsa Cziráky Londe. 2006. Stochastic phonological knowledge: The case of Hungarian vowel harmony. *Phonology* 23: 59–104. doi:10.1017/S0952675706000765.
- Heinz, Jeffrey. 2018. The computational nature of phonological generalizations. In *Phonological typology*, ed. Larry Hyman and Frans Plank, 126–195. De Gruyter Mouton.

Heinz, Jeffrey, Chetan Rawal, and Herbert G. Tanner. 2011. Tier-based strictly local constraints for phonology. Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics 46: 58–64.

Hillenbrand, James, Laura A Getty, Michael J Clark, and Kimberlee Wheeler. 1995. Acoustic characteristics of American English vowels. *Journal of the Accoustical Society of America* 97: 3099–3111.

Kimper, Wendell A. 2016. Asymmetrical generalisation of harmony triggers. In *Proceedings of the 3rd Annual Meeting on Phonology*, ed. Gunnar Ólafur Hansson, Ashley Farris-Trimble, Kevin McMullin, and Douglas Pulleyblank, 1–12. Washington, DC: Linguistic Society of America.

Lai, Regine. 2015. Learnable vs unlearnable vowel harmony patterns. *Linguistic Inquiry* 46: 425–451. doi:10.1162/LING_a_00188.

McMullin, Kevin. 2016. Tier-based locality in long-distance phonotacites: Leanability and typology. PhD Dissertation, University of British Columbia.

McMullin, Kevin, and Gunnar Ólafur Hansson. 2014. Locality in long-distance phonotactics: evidence for modular learning. In *Proceedings of NELS 44*.

Moreton, Elliott. 2008. Analytic bias and phonological typology. *Phonology* 25: 83–127. doi:10.1017/S0952675708001413.

Moreton, Elliott, and Joe Pater. 2012. Structure and substance in artificial-phonology learning, Part II: Substance. Language and Linguistics Compass 6: 702–718. doi:10.1002/lnc3.366.

- Moreton, Elliott, Joe Pater, and Katya Pertsova. 2017. Phonological concept learning. *Cognitive Science* 41: 4–69. doi:10.1111/cogs.12319.
- Pycha, Anne, Pawel Nowak, Eurie Shin, and Ryan Shosted. 2003. Phonological rule-learning and its implications for a theory of vowel harmony. In *West Coast Conference of Formal Linguistics* 22, 101–113.

R Development Core Team, R. 2018. R: A Language and Environment for Statistical Computing. Edited by R Development Core Team. *R Foundation for Statistical Computing*. R Foundation for Statistical Computing. R Foundation for Statistical Computing. doi:10.1007/978-3-540-74686-7.

- Rebrus, Péter, and Miklós Törkenczy. 2015. Monotonicity and the typology of front/back harmony. *Theoretical Linguistics* 41: 1–61.
- Ringen, Catherine O. 1980. A concrete analysis of Hungarian vowel harmony. In *Issues in vowel harmony*., ed. R. M. Vago, 134–154. Amsterdam: John Benjamins.
- Ringen, Catherine O. 1988. Transparency in Hungarian vowel harmony. *Phonology* 5: 327–342. doi:10.1017/S0952675700002335.
- Ringen, Catherine O, and Robert M. Vago. 1998. Hungarian vowel harmony in Optimality Theory. *Phonology* 15: 393–416.
- Smith, Caitlin. 2016. A gestural account of neutral segment asymmetries in vowel harmony. In *Proceedings of the Annual Meeting of Phonology*, ed. Gunnar Ólafur Hansson, Ashley Farris-Trimble, Kevin McMullin, and Douglas Pulleyblank, 3:1–12. Washington, DC: Linguistic Society of America.
- Vago, Robert M. 1976. Theoretical implications of Hungarian vowel harmony. Linguistic Inquiry 7: 243–263.
- Vaux, Bert. 2000. Disharmony and derived transparency in Uyghur vowel harmony. In Proceedings of NELS 30, ed. M. Hirotani, A. Coetzee, N. Hall, and J. Y. Kim, 671–698. Amherst, MA: GLSA.
- White, James Clifford. 2014. Evidence for a learning bias against saltatory phonological alternations. *Cognition* 130: 96–115. doi:10.1016/j.cognition.2013.09.008.
- White, James Clifford, 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 asy

mmetry: Evidence from artificial language learning. In Proceedings of NELS 48.

- Wilson, Colin. 2003. Experimental investigations of phonological naturalness. In West Coast Conference of Formal Linguistics 22, 101–114.
- Wilson, Colin. 2006. Learning phonology with substantive bias: An experimental and computational study of velar palatalization. *Cognitive Science* 30: 945–982.
- Yin, Sora Heng, and James Clifford White. 2018. Neutralization and homophony avoidance in phonological learning. *Cognition* 179: 89–101. doi:10.1016/j.cognition.2018.05.023.