Generalization to Novel Consonants in Artificial Grammar Learning

Sara Finley (sfinley@bcs.rochester.edu)
Department of Brain and Cognitive Sciences, Meliora Hall, University of Rochester
Rochester, NY 14627 USA

Abstract

While theoretical phonologists rely on abstract phonetic features to account for the variety of phonological patterns that exist in the world's languages, it is unclear whether such abstract representations bear psychological reality. Previous research has shown that learners in artificial grammar learning experiments are able to generalize a newly learned phonological pattern to novel segments, suggesting that learners are able to form abstract, feature-based representations. However, conflicting results suggest that this level of abstraction may be restricted to vowels, rather than consonants. The present experiment extends previous findings on generalization to novel segments in vowel harmony to an analogous pattern, consonant harmony. We show that learners fail to generalize to novel consonants in consonant harmony, but succeed at generalization to novel consonants in a general, consonant deletion pattern. Implications for the role of distinctive phonetic features in phonological learning are discussed.

Keywords: statistical learning, consonant harmony, phonology.

Phonetic Features in Phonology

Phonological processes are systematic patterns of sounds that govern the way words and phrases may sound in different languages. Typically, descriptive accounts of these patterns involve abstract, phonetic features to characterize the regularities of the sound patterns. Vowel and consonant harmony is characterized by agreement for a particular vowel/consonant feature value. For example, in back vowel harmony (which occurs in several languages such as Finnish, Hungarian, and Turkish (van der Hulst & van de Weijer, 1995)), if the first vowel of a word bears the feature [back] (e.g., [o, u]), all vowels of the word must also bear the [back] feature. The most concise description of vowel and consonant harmony makes use of distinctive phonetic features such as [back].

There are two ways in which phonological processes can be described without the use of distinctive features. The first is to characterize phonological patterns in terms of the individual segments involved. While this can describe the data, such an analysis misses the generalization that the vowels involved in harmony are not from arbitrary sets, but phonetically similar units. A second option is to characterize the phonological pattern in terms of an exemplar network. This exemplar network would contain a specialized model of similarity that would be used to capture the generalization about the feature involved in the phonological process without specific reference to features (Pierrehumbert, 2001). Such exemplar networks can describe the same phonological processes that make use of features, but are more complex than a feature-based representation because exemplar models require networks of association between all spoken and heard tokens, as well as an algorithm for computing similarity between exemplars.

Previous attempts at providing evidence for the psychological reality of distinctive phonological features have brought mixed results. Using the artificial grammar learning paradigm, researchers trained participants on a novel phonological pattern with a limited sound inventory (e.g., four vowels of a six-vowel inventory). At test, these participants were given items that contain sounds not included in the exposure inventory (e.g., the two vowels held-out of the six-vowel inventory). If learners extend the newly learned phonological pattern to the novel segments, it suggests that they learned the pattern in terms of distinctive phonological features. If they fail to extend the pattern to the novel sounds, it suggests that the learners used some other representation to encode the novel pattern (e.g., based on segment-specific representations). When participants were exposed to a novel phonological pattern involving vowels (e.g., vowel harmony), there was significant generalization to the novel vowels (Chambers, Onishi, & Fisher, 2010; Finley & Badecker, 2009; Skoruppa & Peperkamp, 2011). However, when participants were exposed to a phonological pattern involving generalization to novel consonants, learners failed to generalize (Peperkamp & Dupoux, 2007; Peperkamp, Le Calvez, Nadal, & Dupoux, 2006).

There are several important differences between the studies that did not show generalization to novel segments and the studies that did. Finley and Badecker (2009) and Skoruppa and Peperkamp (2011) used a process that involved vowels, while Peperkamp and Dupoux (2007) trained participants on a process involving intervocalic voicing in obstruent consonants (that applied either only to stops or only to fricatives). It is possible that vowel representations are more amenable to feature-based generalization than consonant representations. There are several reasons why this might be true. The acoustic space of vowels is more variable and continuous than the acoustic space of consonants, which are represented in a more discrete manner. In addition, there is evidence that consonants are processed differently than vowels. Nespor and colleagues have argued that vowels tend to have a more grammatical function, while consonants tend to have a more lexical function (Bonatti, Pena, Nespor, & Mehler, 2005; Toro, Bonatti, Nespor, & Mehler, 2007; Toro, Shukla, Nespor, & Endress, 2008).

Another possibility for the mixed results on generalization to novel segments is that all three studies used relatively different methods for training and testing participants. Skoruppa and Peperkamp (2011) exposed participants to the harmony pattern via listening to a story spoken in a novel accent (Maye, Aslin, & Tanenhaus, 2008). This highly implicit training procedure may have helped learners form abstract representations for the phonological pattern. While Finley and Badecker (2009) and Peperkamp and Dupoux (2007) both used artificial grammar learning paradigms, there are several important differences between the two studies. Specifically, participants in Peperkamp and Dupoux (2007) were trained via a picture-naming task that required participants to map form and meaning, while participants in Finley and Badecker (2009) were trained only on sound forms with no semantic representations. Using purely phonological representations may have made the task slightly easier and thereby made generalization to novel segments easier.

In order to better understand why previous studies have failed to show generalization to novel consonants, we replicated Finley and Badecker (2009) with consonant harmony. If the differences between Finley and Badecker (2009) and Peperkamp and Dupoux (2007) are merely methodological, then participants should generalize to novel consonants in consonant harmony, when trained and tested in a manner identical to Finley and Badecker (2009). If feature-based generalization is based more on the representation of consonants and vowels, and the nature of the specific phonological process, learners will not generalize to novel consonants in consonant harmony.

Experiment 1

Participants

All participants were adult native speakers of English with no previous participation in any experiment involving reduplication. Forty-eight University of Rochester students and affiliates were paid $10 for their participation.

Design

The design of the experiment mirrored Finley and Badecker (2009) as closely as possible. Participants in the critical (trained) conditions were exposed to a consonant harmony pattern. The consonant harmony pattern was based on the feature [+continuant]. This feature distinguishes between stops [d, t, b, p] and fricatives [z, s, v, f]. Participants were exposed to the pattern via pseudo-morphophonological alternations in which a harmonic stem (CVCV) was followed by a harmonic suffixed item (CVCV-bi or CVCV-vi). The suffix was either [bi] or [vi], depending on the nature of the consonants in the stem were stops (triggering [bi]) or fricatives (triggering [vi]).

While consonant harmony based on the continuant feature is not common among consonant harmony languages (Hansson, 2001), the continuant feature was chosen because the English sound inventory has several stop-fricative contrasts, making it relatively easy to design an experiment that tests for generalization to novel consonants.

Participants were divided into three conditions: (Voiceless) Labial Hold-Out, and (Voiceless) Coronal Hold-Out, and a matched Control condition (half of participants were matched to the Coronal Hold-Out condition, and the other half were matched to the Labial Hold-Out condition).

Voiceless Labial Hold-Out Condition Participants in the Labial Hold-Out condition were exposed to 24 pairs of stem, stem+suffix items (repeated five times each), in which the first item contained a CV.CV word, and the second item repeated the first word followed by the [-bi]/[-vi] suffix. All words obeyed the continuant harmony pattern; half the items contained stops only and took the suffix [-bi], while the other half of the items contained only fricatives and took the suffix [-vi]. Training items contained stops and fricatives from the set [s, t, b, v, z, d]. Specifically, the voiceless labial sounds [f/p] were not presented in training.

Following exposure, participants were given a two-alternative forced choice task with 36 items, from three different conditions of test items. 12 of the items were found in the training set (Old Items), 12 items were not found in the training set, but contained the same consonants and vowels as in the training set (New Items), and 12 of the items contained the held-out sounds [f/p] (New Consonant Items). The test items were of the form CVCV-vi vs. CVCV-bi with [-vi] and [-bi] counterbalanced for order of presentation, (e.g., [bobobi] vs. *[bodevi]). Participants were told to select the pair of words that best represented the language they had heard prior to the test.

Voiceless Coronal Hold-Out Condition Participants in the Voiceless Coronal Hold-Out condition were given the same types of training items as participants in the Voiceless Labial Hold-Out Condition, except that the stem items were drawn from the set [p, f, b, v, z, d], and, the voiceless coronal sounds [s/t] were not presented in training.

The test items were similar to those given to the Labial Hold-Out Condition, except that the New Consonant items contained [s/t].

Control Conditions Two Control conditions were created in the same manner as Finley and Badecker (2009), one matched to the Labial Hold-Out condition, and a second matched to the Coronal Hold-Out condition. Participants in the Control condition were given items that would not reflect consonant harmony. The items were a mix of stems that contained harmonic items (those found in the training set for the matched critical condition) and disharmonic items (which were different from the matched critical condition). There were no suffixed items. At test, participants were given the same items as those in the matched critical conditions. While all items are technically ‘new’ (as no suffixed items appeared in training), all classifications are based on the critical conditions.

Materials

A female native English speaker produced all spoken linguistic materials. All items were recorded in a sound-attenuated booth. The speaker had no knowledge of the design or purpose of the experiment. All spoken stimuli contained only CV syllables. Consonants were drawn from the set: /p, t, b, d, s, z, v, z/ and vowels were taken from the set /e, i, o, u/. Care was taken so that all of the stimuli were non-words in English. Examples of training stimuli can be found in Table 1.

Table 1: Training Items.

Voiceless Coronal Hold-Out	Voiceless Labial Hold-Out
vozo, vozovi	vusi, vusivi
sezi, sezivi	sezi, sezivi
didu, didubi	didu, didubi
bupo, bupobi	tute, tutebi

Test stimuli were recorded in the same manner as training stimuli. There were 36 test items, 12 containing pairs of words that appeared in training (Old Items), 12 containing items not heard in training (New Items), and 12 items containing consonant sounds that did not appear in training (New Consonant Items). Items appearing in the New Items were drawn from the same set of consonant and vowels as the training stimuli. Examples of test stimuli are provided in Table 2.

Table 2: Test Items.

Voiceless Coronal Hold-Out	Voiceless Labial Hold-Out
Old Items	
*diduvi vs. didubi	*diduvi vs. didubi
vozovi vs. *vozobi	vusivi vs. *vusibi
New Items	
*dibovi vs. dibobi	sesivi vs. *sesibi
zifuvi vs. zifubi	*botevi vs. botebi
New Consonant Items	
susovi vs. susobi	fefuvi vs. *fefuvi
*tetivi vs. tetibu	*pepivi vs. pepibi

Procedure

All phases of the experiment were run in Psyscope X (Cohen, MacWhinney, Flatt, & Provost, 1993). Participants were given both written and verbal instructions, and were debriefed upon completion of the experiment (which took approximately 20 minutes for participants in the trained conditions, and 10 minutes for participants in the Control condition).

All participants were told that they were to be listening to pairs of words from a language they had never heard before. They were informed that there would be questions about the language following exposure, but that they need not memorize the words they heard. Following exposure, participants were given instructions for the test items. Participants were told that they would hear two words. One word was from the language they had just heard, and the other word was not from the language they had heard; if they believed the first word was from the language, they were instructed to press the 'a' key; if they believed the second word was from the language, they were instructed to press the 'l' key.

Results

Proportion of correct responses for all conditions are given in Figures 1 and 2.

Labial Hold-Out We compared the Labial Hold-Out Training condition (mean = 0.64, *CI* ± 0.084) with the matched Control condition (mean = 0.48, *CI* ± 0.043) via a 2X3 mixed-design ANOVA. There was a significant effect of training ($F(1,22) = 14.30$, $p < 0.01$). There was significant effects of test item ($F(2,44) = 4.69$, $p < 0.05$), which was carried by the fact that correct responses to Old and New items were correct more often than responses to New Consonant test items ($F(1,22) = 6.83$, $p < 0.05$). There was also a significant interaction ($F(2,44) = 4.23$, $p < 0.05$), which was carried by a significant difference between New and New Consonant Items in the Labial Hold-Out condition ($t(11) = 3.32$, $p < 0.01$), but not in the Control condition ($t(11) = 0.38$, $p = 0.71$).

To test for generalization to novel consonants, we compared the New Consonant test items between the Labial Hold-Out (mean = 0.50, *CI* ± 0.13) and the Control (mean = 0.48, *CI* ± 0.082) conditions. There was no significant difference ($t(22) = 0.30$, $p = 0.76$), suggesting that participants failed to generalize to novel consonants.

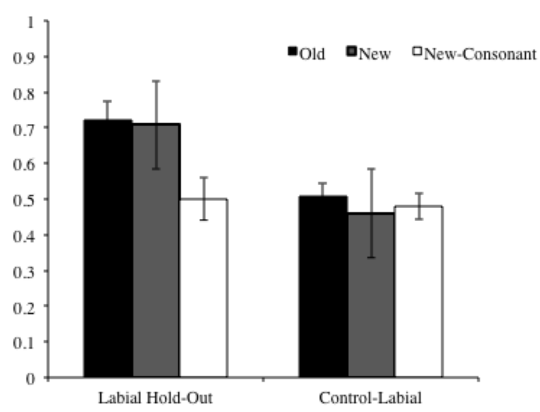

Figure 1: Coronal Hold-Out Results.

Coronal Hold-Out We compared the Coronal Hold-Out Training (mean = 0.67, *CI* ± 0.13) condition with the matched Control condition (mean = 0.42, *CI* ± 0.056) via a 2X3 mixed-design ANOVA. There was a significant effect of training ($F(1,22) = 16.61$, $p<0.01$), suggesting that participants learned the harmony pattern. There were no significant effects of test item (*F*<1). There was a significant interaction ($F(2,44) = 7.35$, $p < 0.01$), which was carried by a significant difference between New and New Consonant Items in the Coronal Hold-Out condition ($t(11) = 2.46$, $p < 0.05$), but not in the Control condition ($t(11) = 1.60$, $p = 0.14$).

To test for generalization to novel consonants, we compared the New Consonant test items between the Coronal Hold-Out (mean = 0.56, $CI \pm 0.12$) and the Control (mean = 0.49, $CI \pm 0.098$) conditions. There was no significant difference ($t(22) = 1.26$, $p = 0.22$), suggesting that participants failed to generalize to novel consonants.

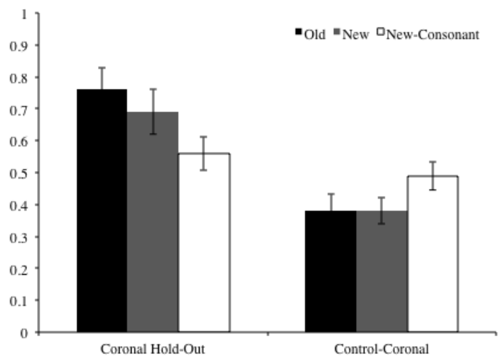

Figure 1: Coronal Hold-Out Results.

Overall Participants were unable to generalize to novel consonants in the consonant harmony pattern, either for novel coronal consonants or novel labial consonants. This failure is not due to a general failure to learn the harmony pattern, as participants were successful in learning the harmony pattern for segments that they were exposed to during training.

Discussion

While Finley and Badecker (2009) showed that adult English-speaking learners of a vowel harmony language generalize to novel segments following brief exposure, participants in the present experiment failed to generalize to novel segments in consonant harmony. There are several possible explanations as to why participants might generalize to novel vowels in vowel harmony but not novel consonants in consonant harmony. First, representations of consonants are very different from representations of vowels. Vowels have a relatively small number of distinctive features (height, rounding, tense, back), while consonants can have several different distinctive features involving both place and manner. The large number of features required to characterize consonants makes it possible to form natural classes of a very small set of consonants. The consonant harmony pattern in the present experiment (as well as the voicing pattern used in Peperkamp and Dupoux, 2007) only involved a select set of stops and fricatives, which is a relatively limited set of segments.

While natural classes typically involve traditional distinctive features, there are rules that apply to sets of consonants that do not form a natural class. This means that it is possible that learners inferred that the consonant harmony pattern applied to a restricted class of segments that excluded the 'held-out' consonants. This is much more likely to occur for consonants than for vowels because the nature of consonants allows for a large number of small, very specific constraints to form a natural class (e.g., coronals and voiced labials). This is more difficult with vowels that are represented in a continuous F1/F2 dimension that do not have many features to begin with. Given that the consonant harmony pattern was highly unnatural, it is likely that learners were more willing to assume a disjoint natural class for the harmony pattern.

Nespor and colleagues have found that speakers process vowels differently from consonants (Bonatti, et al., 2005; Toro, et al., 2007; Toro, et al., 2008). They hypothesize that consonants are processed in terms of lexical content, while vowels are processed in terms of their grammatical function. If learners are biased to process phonological patterns involving consonants in terms of lexical processing rather than grammatical function, learners may be more likely to infer a more select, strict classification of a consonant harmony process than a vowel harmony process.

Given that learners were unable to generalize to novel consonants in consonant harmony, there is a question of whether learners can generalize to novel consonants after learning any novel pattern involving consonants, even a highly general one. If learners are given a process that applies to all consonants, such as an across-the-board deletion process, then it is likely that learners will generalize to the novel consonant.

In Experiment 2, we exposed learners to a highly general phonological process that does not require representation of specific feature values (deletion). In a consonant deletion rule, a constraint against two adjacent consonants causes one of the consonants (typically the first one) to delete (Wilson, 2001). Because the constraint is general against all consonants, and does not involve specific features or feature pairs, it is more likely that learners will form a general rule that applies to all consonants. In Experiment 2, we test for generalization to novel consonants in consonant deletion.

Experiment 2

Participants

All participants were adult native speakers of English with no previous participation in any experiment involving consonant deletion. Twenty-six University of Rochester undergraduate students and affiliates were paid either $10 or $5 for their participation (participants in the No-Training Control condition were paid $5).

Design

The design of Experiment 2 mirrored Experiment 1 in that participants were exposed to a phonological pattern based on pseudo morpho-phonological alternations. In Experiment 2, the phonological pattern was a consonant deletion pattern that was presented to participants in sets of three words: A, B, and AB (a combination of A and B). The words in A, and B were of the form CVC and the combined AB form deleted the final consonant of the A word (e.g., biv mop, bimop). The deleted consonant was always a fricative or a nasal from the set /s, z, v, z, f, m, n/, and never a stop consonant.

Following exposure, participants were given a two-alternative forced choice task with 30 items, of the same types as Experiment 1: Old Items, New Items, and New Consonant Items. In New Consonant Items, the deleted item was always a stop consonant /p, t, k, b, d, g/. Participants were told to select the pair of words that best represented the language they had heard prior to the test.

Materials

Spoken materials were produced by a female native English speaker in a sound-attenuated booth. The speaker had no knowledge of the design or purpose of the experiment. Consonants were drawn from the set: /p, t, k, b, d, g, s, z, v, z, f, m, n/ and vowels were taken from the set /a, e, i, o, u/. Care was taken so that all of the stimuli were non-words in English. Examples of training stimuli can be found in Table 3, below.

Table 3: Experiment 2 Training Items.

A (CVC)	B (CVC)	AB (CVCVC)
pim	bop	pibop
fev	sof	fesof
buf	ven	buven

Procedure

The procedure of Experiment 2 was identical to Experiment 1, except that participants in Experiment 2 were told that they were listening to words in sets of three, and the third word would be a combination of the first two words (just as ‘tooth’ and ‘brush’ combine to form ‘toothbrush’).

Results

Proportion of correct responses for all conditions are given in Figure 3.

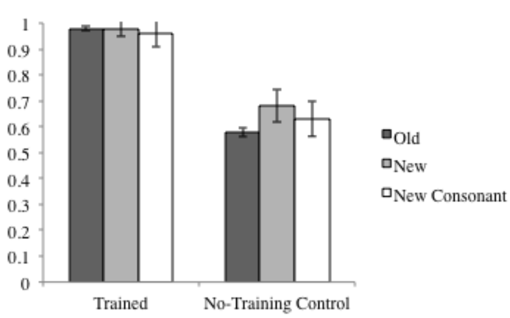

Figure 3: Results.

We compared the Trained Condition (mean = 0.98, *CI* ± 0.012) to the No-Training control condition (mean = 0.63, *CI* ± 0.11) via a 2X3 mixed-design ANOVA. There was a significant effect of training ($F(1,24) = 44.98$, $p < 0.0001$), suggesting that participants learned the deletion pattern. There was no significant effect of test item ($F(2,48) = 1.44$, $p = 0.25$), and no interaction ($F(2,48) = 2.00$, $p = 0.14$).

We also compared the Training (mean = 0.96, *CI* ± 0.039) and the Control (mean = 0.63, *CI* ± 0.15) conditions specifically for the New Consonant Items, in order to assess generalization to novel consonants. There was a significant difference between the Control and the Trained conditions ($t(24) = 4.74$, $p < 0.0001$), suggesting generalization to the novel consonants.

Discussion

Participants were able to generalize to novel consonants in the consonant deletion pattern. This suggests that generalization to novel consonants can be achieved if the pattern in question is highly general.

General Discussion

This paper explored the situations in which learners are able to generalize to novel consonants in an artificial grammar learning experiment. Previous artificial grammar learning experiments have shown that adult learners can generalize to novel vowels (Chambers, et al., 2010; Finley & Badecker, 2009; Skoruppa & Peperkamp, 2011), but not novel consonants (Peperkamp & Dupoux, 2007; Peperkamp, et al., 2006). Experiment 1 ruled out the possibility that the difference between vowels and consonants was purely methodological. Using an identical paradigm to Finley and Badecker (2009), we exposed learners to a consonant harmony pattern, and then tested generalization of that novel consonant harmony pattern; learners did not extend the harmony pattern to novel consonants. The methods for consonant harmony learning were identical to those used in previous vowel harmony learning experiments, and the consonant harmony pattern was parallel to the vowel harmony learning paradigm. Therefore, one must conclude that the representations of the consonant harmony pattern were subject to different constraints than vowel harmony. Because vowel harmony typically involves all the vowels in the inventory (excluding only a small subset of vowels), but consonant harmony typically only involves a small subset of consonants (e.g., sibilant consonants), and excludes a large number of consonants in the language, consonant harmony patterns are much more likely to be restricted to a small, relatively idiosyncratic set of consonants. This may lead learners to form a restrictive set of consonants that follow the consonant harmony pattern, leading to a failure to generalize to novel consonants.

If generalization to novel consonants depends on a phonological pattern that applies to a large set of vowels/consonants, consonant harmony may not be the most ideal case. In Experiment 2, we tested learners’ ability to generalize novel consonants to a highly general deletion pattern. In this pattern, the first consonant of a pair of adjacent consonants was deleted. During training, the deleted item was always a fricative. At test, learners generalized to stops, which had not participated in the deletion process. Participants generalized to the novel stops, suggesting that the learners had formed a general representation of the deletion process.

The results of the present set of experiments demonstrate that the level of generality of representation depends largely on the type of process in question. A consonant harmony pattern that requires agreement between pairs of stops and fricatives is highly specific, and therefore subject to specific interpretation. Deletion, however, is highly general, applying to all consonants, and is therefore subject to a highly general interpretation.

The high degree of specificity found in patterns in consonant harmony is consistent with the theories of vowels and consonants put forth by Nespor and colleagues discussed above (Bonatti, et al., 2005; Toro, et al., 2007; Toro, et al., 2008). One reason that processes involving consonants tend to carry more lexical information is that there is a much higher degree of variation among consonants than vowels. This means that patterns involving consonants should be more likely to be subject to specific representations compared to vowels, which vary on a smaller number of dimensions compared to consonants.

Conclusions

While previous experiments investigating feature-based generalization in consonants have failed to show generalization to novel consonants, the present study demonstrated for the first time that learners can generalize to novel consonants when the phonological pattern is highly general. Learners of a novel consonant deletion pattern were able to generalize to consonants that had not previously undergone deletion. However, when exposed to a highly specific consonant harmony pattern, learners failed to generalize to novel consonants. This suggests that learners are able to form highly general representations for novel phonological patterns when these patterns apply to a wide range of consonants/vowels.

Acknowledgments

The author would like to thank three anonymous reviewers, Patricia Reeder, Elissa Newport, Neil Bardhan, and members of the Aslin-Newport Lab. In addition, we would like to thank Carrie Miller, Kelly Johnston, Emily Kasman, and Anna States for help in stimuli creation and data collection. I assume all responsibility for any errors. Funding was provided by NIH Grants DC00167 and T32DC000035.

References

Bonatti, L. L., Pena, M., Nespor, M., & Mehler, J. (2005). Linguistic constraints on statistical computations: The role of consonants and vowels in continuous speech. *Psychological Science, 16*, 451-459.

Chambers, K., Onishi, K. H., & Fisher, C. (2010). A vowel is a vowel: Generalization to newly learned phonotactic constraints to new contexts. *Journal of Experimental Psychology: Learning, Memory and Cognition, 36*(3), 821-828.

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.

Finley, S., & Badecker, W. (2009). Artificial grammar learning, and feature-based generalization. *Journal of Memory and Language, 61*, 423-437.

Hansson, G. O. (2001). *Theoretical and typological issues in consonant harmony.* Ph.D. dissertation, University of California, Berkeley, Berkeley.

Maye, J., Aslin, R. N., & Tanenhaus, M. K. (2008). The weckud wetch of the wast: Lexical adaptation to a novel accent. *Cognitive Science, 32*, 543-562.

Peperkamp, S., & Dupoux, E. (2007). Learning the mapping from surface to underlying representations in artificial language. In J. Cole & J. I. Hualde (Eds.), *Laboratory Phonology 9* (pp. 315-338).

Peperkamp, S., Le Calvez, R., Nadal, J.-P., & Dupoux, E. (2006). The acquisition of allophonic rules: Statistical learning with linguistic constraints. *Cognition*, B31-B41.

Pierrehumbert, J. (2001). Exemplar dynamics: Word frequency, lenition and contrast. In J. L. Bybee & P. Hopper (Eds.), *Frequency effects and emergent grammar* (pp. 137-157). Amsterdam: John Benjamins.

Skoruppa, K., & Peperkamp, S. (2011). Adaptation to novel accents: Feature-based learning in context-sensitive phonological regularities. *Cognitive Science, 35*(2), 348-366.

Toro, J. M., Bonatti, L., Nespor, M., & Mehler, J. (2007). Finding words and rules in a speech stream: Functional differences between consonants and vowels. *Psychological Science, 19*(2), 137-144.

Toro, J. M., Shukla, M., Nespor, M., & Endress, A. (2008). The quest for generalizations over consonants: Asymmetries between consonants and vowels are not the by-product of acoustic differences. *Attention Perception and Psychophysics, 70*(8), 1515-1525.

van der Hulst, H., & van de Weijer, J. (1995). Vowel harmony. In J. Goldsmith (Ed.), *The handbook of phonological theory* (pp. 495-531). Oxford: Blackwell.

Wilson, C. (2001). Consonant cluster neutralisation and targeted constraints. *Phonology, 18*, 147-197.