Published in final edited form as:

J Mem Lang. 2011 July 1; 65(1): 74–83. doi:10.1016/j.jml.2011.02.006.

The privileged status of locality in consonant harmony

Sara Finley

Department of Brain and Cognitive Sciences, University of Rochester, Meliora Hall, Rochester, NY 14627

Sara Finley: sfinley@bcs.rochester.edu

Abstract

While the vast majority of linguistic processes apply locally, consonant harmony appears to be an exception. In this phonological process, consonants share the same value of a phonological feature, such as secondary place of articulation. In sibilant harmony, [s] and [ʃ] ('sh') alternate such that if a word contains the sound [ʃ], all [s] sounds become [ʃ]. This can apply locally as a first-order or non-locally as a second-order pattern. In the first-order case, no consonants intervene between the two sibilants (e.g., [pisasu], [piʃaʃu]). In second-order case, a consonant may intervene (e.g., [sipasu], [ʃipaʃu]). The fact that there are languages that allow second-order non-local agreement of consonant features has led some to question whether locality constraints apply to consonant harmony. This paper presents the results from two artificial grammar learning experiments that demonstrate the privileged role of locality constraints, even in patterns that allow second-order non-local interactions. In Experiment 1, we show that learners do not extend first-order non-local relationships in consonant harmony to second-order nonlocal relationships. In Experiment 2, we show that learners will extend a consonant harmony pattern with second-order long distance relationships to a consonant harmony with first-order long distance relationships. Because second-order non-local application implies first-order non-local application, but first-order non-local application does not imply second-order non-local application, we establish that local constraints are privileged even in consonant harmony.

Locality constraints have played an important role in characterizing linguistic processes. In syntax, morphology and phonology, local processes tend to be privileged over non-local processes (Chomsky, 1981; Culicover & Wilkins, 1984). In phonology, the bias for local processes is based on phonetic restrictions. Many phonological processes are a result of coarticulation, which is the process of overlapping gestures of sounds (e.g., the [i] in [bin] carries some of the nasal properties of the final [n] consonant). Because coarticulation is likely to decrease over distance, phonologization of coarticulation is likely to be based on adjacent sounds (e.g., nasalization of the [i] but not the [b] in [bin], as coarticulation will not carry all the way to the first consonant) (Hyman, 1976; Ohala, 1994).

In addition to phonetic principles, locality restrictions in language may have non-linguistic roots. For example, constraints on perception and memory may serve as a source for constraints on locality in language. General cognitive processes such as categorization and memory are often based on locality. Two objects are more likely to be placed in the same category if they are grouped closer together in space than a third object that is farther away. In memory, people use adjacency as part of memory aids, such as 'chunking' (e.g., phone

numbers are chunked in terms of adjacent digits (310-555-1234) rather than in terms of odd and even numbers, which may not be adjacent).

The combination of phonetic principles and general cognitive biases towards locality has led to the assumption that phonological patterns must be governed by principles of locality. Such principles of locality are able to account for the majority of phonological processes, but there are some processes that appear to allow (if not prefer) non-local interactions. Consonant harmony is a specific case of a linguistic pattern that allows non-local interactions. In this phonological process, consonants must share the same value for a particular phonological feature, such as secondary place of articulation.

Because a vowel may intervene between the two agreeing consonants, consonant harmony is considered a non-local process; agreeing consonants need not be adjacent. There are two levels of non-locality in consonant harmony. In first-order non-locality, the two consonants involved in harmony may be separated by vowels, but may not have any intervening consonants. For example, in nasal harmony in Ndonga, [l] becomes [n] after a nasal (e.g., [m] and [n]) sound (e.g., [k~~u~~n-in-a] ‘sow for’). However, nasal harmony fails to apply if a consonant intervenes between the nasal [n] and the liquid [l], as in [nik-il-a] ‘season for’ (Rose & Walker, 2004). In this case of first-order non-locality, vowels and consonants are separate from each other; vowels do not block harmony.

Second-order non-locality occurs when vowels as well as irrelevant consonants may intervene between the two agreeing consonants. Such second-order patterns occur in Navaho consonant harmony, which requires agreement between sibilant fricatives [s] and [ʃ] (‘sh’). In sibilant harmony, [s] and [ʃ] alternate such that if a word contains the sound [ʃ], all preceding (or following) [s] sounds become [ʃ]. In Navaho, the perfective prefix alternates between [si-] and [ʃi-] (Hansson, 2001; Heinz, 2010; McDonough, 1990; Rose & Walker, 2004; Sapir & Hojier, 1967). The prefix [si-] appears if the stem contains the alveolar fricatives or affricates /s, z, ts, tsʰ, dz/, as in [site:z] ‘my car’, while [ʃi-] appears if the stem contains the postalveolar fricatives or affricates /ʃ, tʃ (‘ch’), ʒ, ʒz/, as in [ʃitʃii:h] ‘my nose’. In second-order non-locality, non-relevant consonants may intervene between the two consonants participating in harmony (e.g., [ʃite: z] ‘we are lying’). Languages with consonant harmony allow either first-order non-locality, or both first and second-order non-locality, but never require second-order non-locality without first-order non-locality. Second-order non-locality therefore implies first-order locality, but first-order locality need not imply second-order locality.

This implication can be traced to phonetic principles of coarticulation. Phonological assimilation tends to be grounded in coarticulation, which is strongest between adjacent segments (Öhman, 1966; Vilain, Abry, & Badin, 2000). However, coarticulation can pass through intervening segments. Öhman (1966) showed that speakers co-articulate vowel features even when a consonant intervenes between the two vowels. Coarticulatory effects are the greatest at the adjacent segment, and gradually reduce as the sounds move further from the source. This means that (i) coarticulation is biased to be local and that (ii) long distance coarticulation implies local coarticulation, as local coarticulation serves as an intermediate step for non-local effects.

The idea that non-local phonological processes must be rooted in coarticulation has been formalized in phonology as a theory of *strict locality* (Gafos, 1999). Strict locality proposes that non-local phonological processes apply at the phonetic level, even if application of the process only changes subtle properties of the intervening segment, and does not change the phonemic category of segment (NiChiosain & Padgett, 2001). According to the strict locality hypothesis, spreading between non-adjacent segments involves gestures that carry

the harmonic feature between the non-adjacent items. In nonlocal harmony, spreading applies from the source to the target through intervening segments. However, the strength of the coarticulation in the intervening segments is not strong enough to warrant a categorical change in the vowel features. For this reason, intervening segments are interpreted as non-undergoing at the phonological level. Because speakers produce coarticulation, but do not consciously report hearing the effects of the coarticulation, harmonic gestures are best (and possibly only) detected through sophisticated instrumentation.

Evidence for the strict locality hypothesis comes largely from vowel harmony, a phonological pattern similar to consonant harmony, but applies to featural agreement among vowels. In vowel harmony, second-order restrictions emerge when the set of vowels that fail to spread their feature value, but the following vowels take the feature specification of the closest participating vowel. These vowels are said to be ‘transparent’ because active vowels will spread right through them (as in Hungarian, in which a back vowel will ‘spread’ its backness feature to another vowel when the front vowels /i/ or /e/ intervene (e.g., [radir-nak] ‘eraser’ DATIVE’ (Ringen & Vago, 1998))). Transparent vowels in vowel harmony are the only case of second-order non-locality in vowel harmony, and are highly restricted. Most languages with vowel harmony do not have transparent vowels. The restrictions on second-order locality in vowel harmony have been instantiated in learning. Finley (in revisions) demonstrated that adult learners of a vowel harmony pattern were able to learn a vowel harmony pattern with an intervening vowel that either obeyed first-order locality or second-order locality (transparent vowels). Participants were only able to learn the second-order locality pattern after increases in exposure time and training items. These results suggest that learners are biased towards first-order locality in vowel harmony.

Benus and Gafos (2007) demonstrated that transparent vowels in Hungarian undergo small but significant gestural coarticulation. However, this coarticulation does not cause enough change for listeners to perceive the vowel undergoing coarticulation as having changed its feature value. Benus and Gafos (2007) suggest that all vowels undergo spreading, even if they are not categorically classified as such. Thus, nonlocality in vowel harmony can be characterized in terms of a single phonetic gesture. This suggests that vowel harmony can apply as a single spreading operation. The result is that vowel harmony only allows second-order dependencies only in highly specific instances. As mentioned above, coarticulatory effects are weaker for consonants than for vowels, meaning that it is less likely that second order consonant harmony patterns will satisfy the coarticulation requirements for a strict locality analysis. For this reason, Hansson (2001) argues that strict locality may be less likely to apply to consonant harmony because the harmonic features in consonant harmony (e.g., sibilant) are less likely to spread via coarticulation.

To deal with the differences in coarticulation between vowels and consonants, Gafos (1998) proposes that apparent violations of strict locality in consonant harmony are due to the fact that such cases are not ‘true’ consonant harmony, but a *feature-copying* process (Gafos, 1998; Nevins, 2010; Rose & Walker, 2004). While Gafos argues that such feature-copying cases of harmony are not technically ‘harmony’, we follow Hansson (2001) and assume for the purposes of this paper that all cases of agreement between consonants (local or at a distance) should be classified as consonant harmony.

Feature-copying proposals of consonant harmony are not subject to the same locality restrictions as strict locality. Rose and Walker’s (2004) feature-copying proposal allows consonant harmony to apply between an unlimited set of intervening segments through correspondence between the two agreeing segments (e.g., $s_1 X s_1$). When two segments are in correspondence, they must share the same value of the harmonic feature, but these corresponding segments need not be adjacent. Heinz (2010) proposes a similar mechanism

using precedence relations. A precedence relation is a pair of segments in which the first segment precedes the second (even if another segment intervenes). For example, the word [sitas] contains the following precedence relations between consonants: [st, ss, ts]. When statistics are calculated over the precedence relations, a pattern will emerge that [ss] is a common precedence relation, but [sf] and [fs] are not found. This establishes the restriction on agreement between sibilant consonants. Because the precedence relationships are stated independently of intervening segments, any number segments may intervene.

Feature copying proposals allow for long distance processes, while strict locality requires adjacency on the surface. Another proposal for accounting for locality violations in consonant harmony is through hierarchical representations. In the *tiers approach* to consonant harmony (also referred to as autosegmental phonology (Goldsmith, 1979)), processes that appear non-local on the surface (e.g., harmony between two consonants separated by a vowel) can receive a local analysis by representing vowels and consonants separately on different tiers (as in Table 1). Harmony is local in this case, but only with respect to the representations proposed by the theorist. In strict locality, local relationships must persist at the surface, beyond representational assumptions. It is not enough to propose separate tiers for vowels and consonants, the harmony process must apply to the two consonants in addition to the intervening vowel.

The tiers approach to consonant harmony is grounded in theories of phonetic representations. Keating (1988) argues that phonetic underspecification will determine whether a process will apply locally or at a distance. Phonetic underspecification results when a sound is not marked (either in the phonology or in the lexicon) for a particular phonetic implementation. When a sound is phonetically unspecified, the speaker may simply carry over the articulatory gesture from an adjacent sound (this process is referred to as phonetic interpolation), pronouncing a specified value from another segment. This produces the effects of coarticulation and spreading. When a series of unspecified segments occurs, the segments will take on the co-articulated feature value at a distance, producing non-local spreading. This proposal can be adapted to conform to tier-based proposals for consonant harmony. For example, providing a separate representational space for all sibilant consonants allows segments specified as sibilants to interact with each other as if they were adjacent, even if non-sibilant segments intervene. If the sibilant feature is given its own tier, segments that have no sibilant feature (e.g., voiceless stops such as /t/) will have no representation on the sibilant tier. Thus, two sibilant sounds will appear adjacent on the sibilant tier, and will be able to undergo spreading constraints like consonant harmony despite having intervening consonants on the surface form.

The present study tests the privileged nature of first-order non-local interactions in consonant harmony using two adult artificial grammar learning experiments. In an artificial grammar learning paradigm with adult speakers, adults are given a sample of a miniature, experimenter-designed language that conforms to some linguistic pattern(s) (consonant harmony in the present case). Following exposure, participants are given a test of words in the language, in order to assess whether learners inferred the pattern they were trained on. It is also possible to test participants on novel items that require generalization beyond the initial training set. This allows the experimenter to test the level of representation that learners infer from the exposure materials. The artificial grammar learning paradigm is an ideal method for testing the role of locality in long-distance consonant harmony processes. Unlike natural language learning settings, it is possible to explore two different types of language (first and second order non-locality) with minimal differences. It is also possible to test inferences to novel material that might not otherwise be possible in a natural setting. Further, given that the study of linguistic universals is fraught with confounds, finding

support for universals in adult learners provides strong support for previously asserted universal principles (Nevins, 2009).

Previous research using the artificial grammar learning paradigm has shown that adult learners are sensitive to implicational universals (Wilson, 2006). An implicational universal is a tendency found among the majority of languages in the world (i.e., a cross-linguistic tendency) in which if a language has particular property X, it implies that the language will also have an additional property Y, but having the property Y does not imply having property X. Wilson exposed participants to a velar palatalization, phonological pattern in which a velar consonant (/k, g/) becomes more front (and therefore palatal (e.g., [ʃ] 'sh') before a front vowel. For example, when the [k] in [ki] becomes palatalized, it is pronounced as [tʃi] 'chee'. Velar palatalization follows an implicational universal such that if mid vowels trigger palatalization, high vowels must also trigger palatalization, but not vice versa. This is due to the fact that high vowels are phonetically more likely to trigger fronting than mid vowels. In Wilson's experiment, learners who were exposed to a mid vowel trigger generalized to a high vowel trigger, but learners who were exposed to a high vowel trigger failed to generalize to the mid vowel trigger. These results support the implicational universal of mid and high vowel triggers for fronting. In the case of consonant harmony, there is an implicational universal in nonlocal patterns. If a consonant harmony pattern allows non-local spreading, it will also allow local spreading, but a local pattern may not necessarily entail a non-local pattern.

Previous research using artificial grammar learning to study vowel harmony (Finley & Badecker, 2008, 2009a, 2009b, in press; Pycha, Nowak, Shin, & Shosted, 2003) and consonant harmony (Wilson, 2003) has shown that adult learners can acquire phonological agreement patterns with relatively short training. In these experiments, learners were exposed to pseudo-morphophonological patterns in which allomorphs depended on the harmonic feature of the stem. Participants were more likely to learn harmony patterns that were natural over unnatural patterns. Natural patterns are both non-arbitrary (phonetically grounded) and follow cross-linguistic tendencies. For example, Finley and Badecker (2009a) showed that learners are able to learn a back/round vowel harmony pattern, and extend that pattern to vowels that did not appear in the training set. However, this generalization only occurred when the novel vowels had the features necessary to trigger round harmony. These results suggest that learners form rules using features and natural classes, but are sensitive to the representations required for participation in phonological processes.

The experiments presented in this paper test two hypotheses related to the privileged status of locality in consonant harmony. Experiment 1 tests the hypothesis that exposure to a local pattern should not imply a non-local pattern. Experiment 2 tests the hypothesis that exposure to a non-local pattern should imply a local pattern

EXPERIMENT 1

Experiment 1 exposes adult (English-speaking) learners to a sibilant consonant harmony language that involves a first-order pattern. Participants are tested on whether the local pattern implies a second-order pattern.

Method

Participants—All participants were adult native English speakers with no knowledge of a consonant harmony language, and had not previously participated in a consonant-harmony learning experiment. Twenty University of Rochester undergraduate students and affiliates and were paid \$10 for their participation.

Design—Participants were divided into a control condition and a trained condition. The participants in the trained condition received items that reflected a first-order consonant harmony pattern. The trained participants were compared to participants in the control condition (described in detail below), who received no information regarding the harmony pattern. The test contained items heard in the training set (Old), new first-order items (New), and second-order items (Second-Order).

Materials

Participants in the trained condition were exposed to a first-order consonant harmony language in which a stem was followed by a suffix that alternated between [-su] and [-ʃu]; [-su] surfaced when the preceding consonant was [s] and [ʃ] surfaced when the preceding consonant was [ʃ]. Learners were exposed to stems (of the form CVsV or CVʃV) immediately followed by the ‘suffixed’ counterpart (e.g., [bise bisesu]). All items contained either [s] or [ʃ].

We did not include stem items containing only non-sibilant consonant. Such forms would trigger a ‘default’ suffix forms (this is generally [s] across natural languages). If such forms were included, the total number of [-su] items would be greater than [-ʃu] items. Because we wanted to have an equal number of items taking [-su] and [-ʃu], neutral items were not included. Further, such items were not informative to the harmony pattern, and therefore did not provide useful learning information. It is unclear whether such items would make learning more difficult.

The second consonant in each stem was either [s] or [ʃ], and first consonant was a stop [p, t, k, b, d, g]. Vowels were drawn from the set /a, i, e, o, u/. There were 24 stem- suffix pairs repeated in a random order 5 times each. Examples of training stimuli can be found in Table 2.

A control condition was created to assess biases in the stimuli and/or prior to the experiment. Participants in the control condition were exposed to words from the artificial language, but the words were stems only, and thus provided no evidence for the harmony pattern. Participants in the control condition were given the same test items as participants in the critical condition. The purpose of the control condition was to assess the role of training in the responses of the critical condition. If learners responded based on pre-existing biases towards harmonic/disharmonic sequences, unrelated acoustic differences in the harmonic and disharmonic test items, or based on the exposure to the stem forms (which contained no harmonic information), the results of the trained condition should be identical to those of the control. The control condition therefore serves as a ‘sanity check’, to ensure that any results are due to learning and not unforeseen, unrelated factors. Participants in the control condition were exposed to the same CVsV and CVʃV stems as those heard in the trained condition. However, participants did not hear any suffixed forms, which gave them no access to a harmony pattern. Giving participants some training makes it possible to give the same instructions (training and test), as the trained condition (as opposed to a ‘no training’ condition¹).

Following training, participants were given a two-alternative forced-choice test in which participants chose between a harmonic and a disharmonic word that differed only between the two suffix allomorphs [-su] and [-ʃu] (e.g., [*² bisiʃu vs. bisisu]). Test items included Old Items (from the training set), New Items (not in the training set, but were of the same

¹Previous research (e.g., Finley and Badecker (2009a)) found no differences between ‘no-training’ control conditions and stem-only control conditions.

²The ‘*’ indicates the disharmonic item.

CVsV and CVʃ V forms seen at training³), and Second-Order Items, in which a non-sibilant stop consonant intervened between the two harmonic sibilant items Second- Order Test Items required second-order non-local constraints, and were of the form *[VCVsu vs. ʃVCVʃu]. There were equal proportions of each type of test item (12 items for each test condition).

All stimuli were naturally produced, and recorded in a soundproof booth by an adult female native English speaker. While the volunteer was aware that the stimuli would be used in an artificial grammar learning experiment, she was blind to the design or purpose of the study. Examples of test stimuli can be found in Table 3. The full stimulus list for Experiment 1 can be found in Appendix 1.

If participants learned the general harmony pattern they should select the correct allomorph at a rate significantly greater than chance (50%). If participants inferred a nonlocal harmony rule from the local rule, they should extend the harmony pattern to the non-local Test Items.

Procedure

All phases of the experiment were presented using PsyScopeX (Cohen, MacWhinney, Flatt, & Provost, 1993). Participants were given written and verbal instructions, and were told that they would be listening to words from a language they never heard before. Their task was to listen to the way the novel language sounded, but that they need not memorize the forms. Both the exposure and test phases were presented over headphones (no visual stimuli were provided). The training was followed by a forced-choice test with 36 suffixed items, one harmonic and one disharmonic. Participants were told that one item was from the language and one item was not from the language, and that it was their job to select the word from the language. The task was to press the 'a' key if the first item was more likely to come from the language, and the 'l' key if the second item was more likely to come from the language. Participants were told to respond as quickly and accurately as possible. Participants were given a debriefing statement upon completion of the experiment (which took approximately 15 minutes).

Results

If learners treat local harmony patterns in a privileged manner, participants will learn the local harmony pattern, but fail to generalize to non-local items. We thus expect learners in the trained condition to select the harmonic option more often than in the control condition for Old and New Items, but not for Second-Order Items.

Proportions of harmonic responses were recorded for each subject in the trained and the control conditions, and can be found in Figure 1. Participants in the Control condition were compared to participants in the trained condition via a mixed design ANOVA with alpha set at 0.05. The between-subjects factor was Training, with two levels in the ANOVA. Test Items (Old Stems, New Stems, New Suffix) was a within-subjects factor nested under the between-subjects factor Training. All conditions involved between-item comparisons.

There was a significant effect of Training ($F(1,18) = 6.69$; $p < 0.05$), in that participants in the trained condition were more likely to choose the harmonic option than participants in the control condition (mean = 0.80 vs. 0.47; $CI = \pm 0.13$). There was an effect of Test Item ($F(2, 36) = 4.62$, $p < 0.05$), reflecting greater harmonic responses to Old and New Items compared to Second-Order Items ($F(1, 18) = 5.25$, $p < 0.05$). There was also a significant interaction

³Note that the distinction between 'Old' and 'New' items is only relevant to the trained condition. All items are 'new' to the control condition. However, participants in the Control condition had heard the stems in the Old items (but not as suffixed, harmonic forms).

between Training and Test ($F(2, 36) = 5.78, p < 0.01$), which was carried by the fact that there was a significant difference between Second-Order and New and New Items for the trained condition ($F(1, 9) = 11.60, p < 0.01$), but not the control condition ($F < 1$). This suggests that participants learned the first-order non-local harmony rule but failed to extend the pattern to the second-order non-local cases.

To directly test whether participants generalized to non-local patterns, we performed a t-test between trained and control conditions for the Second-Order Items. There was no significant difference between the trained condition and the control condition ($t(18) = 0.20, p = 0.85$), further suggesting that participants failed to extend the first-order harmony pattern to the second-order cases.

In addition to comparing the trained condition to the control condition, we performed one-sample t-tests to compare each test condition to 50% chance. In the trained condition, only Old ($t(9) = 5.48, p < 0.0001$) and New ($t(9) = 3.64, p < 0.001$) Items were above chance; Second-Order Items were not ($t(9) = 0.08, p = 0.95$). None of the test items in the Control condition were significantly different from chance: Old ($t(9) = 0.57, p = 0.581$) and New ($t(9) = .36, p = 0.73$), and Second-Order ($t(9) = 0.22, p = 0.83$).

Of the 10 participants in the trained condition, only 3 selected the harmonic option in the Second-Order Items at a rate greater than 50%. This suggests that while the majority of participants inferred a first-order harmony pattern, some participants may have inferred a second-order non-local pattern, suggesting that the bias towards local patterns may be probabilistic. We will return to the implications of this in the General Discussion.

Discussion

Participants in Experiment 1 learned the first-order consonant harmony pattern, but failed to infer a second-order non-local harmony pattern. This supports the hypothesis that first-order consonant harmony patterns do not imply second-order patterns. It further suggests that first-order instances of consonant harmony may be privileged.

It is possible that participants in Experiment 1 simply learned a specific rule in which the last two consonants of a word are subject to sibilant harmony, but no other consonants are subject to this constraint. A specific rule of this type would make it possible for participants to choose the harmonic items in the New Test Items, but would be at chance for the Second-Order Test Items, in which the relevant consonant was not the last consonant of the stem.

If learners are biased towards such types of specific patterns, learners should fail to generalize from a second-order case to a first order case. Trained on a second-order case, learners may infer a rule in which first and last consonants agree. However, if learners actually infer a second-order long distance pattern, which implies the presence of a first-order pattern, one should expect generalization to first-order instances of harmony after being trained only on the second-order pattern. This is tested in Experiment 2.

EXPERIMENT 2

Experiment 2 tests the second part of the hypothesis that first-order patterns in consonant harmony are privileged by exposing learners to a second-order non-local pattern and testing for generalization to a first-order pattern. If first-order patterns show a privileged status, a second-order non-local harmony pattern should imply a first-order pattern, particularly in light of the results of Experiment 1, in which first-order patterns did not imply second-order non-local patterns.

Method

Participants—All participants were adult native English speakers with no knowledge of a consonant harmony language, and had not previously participated in a consonant-harmony learning experiment. Twenty University of Rochester undergraduate students and affiliates and were paid \$10 for participation.

Design—The design of Experiment 2 was identical to that of Experiment 1, except that learners were exposed to a second-order non-local harmony pattern in training, and were tested to generalization to the local pattern at test.

Materials

The materials were recorded in the same manner as in Experiment 1. Examples of training and test stimuli can be found in Tables 1 and 2, and the full stimulus list can be found in Appendix 2.

Procedure

The procedure was identical to that of Experiment 1.

Results

If learners treat local harmony patterns in a privileged manner, they will generalize from second-order items to first-order non-local items. We thus expect learners in the trained condition to select the harmonic option more often than in the control condition Old, New and First-Order Test Items.

Proportions of harmonic responses were recorded for each subject in the trained and the control conditions, and can be found in Figure 2. Participants in the Control condition were compared to participants in the trained condition via a mixed design ANOVA with alpha set at 0.05. The between-subjects factor was Training, with two levels in each ANOVA. Test Items (Old Stems, New Stems, Local) was a within-subjects factor nested under the between-subjects factor Training. All conditions involved between-item comparisons.

There was a significant effect of Training ($F(1,18) = 29.21, p < 0.001$), in that participants in the trained condition were more likely to choose the harmonic option than participants in the control condition (mean = 0.80 vs. 0.41; $CI = \pm 0.11$). There was no effect of Test Item ($F < 1$) and no interaction ($F < 1$), suggesting that responses to all three types of test items were equivalent.

To directly test whether participants generalized to non-local patterns, we performed a t-test between trained and control conditions for the First-Order items. There was a significant difference ($t(18) = 3.59, p < 0.01$), suggesting that learners generalized from a non-local pattern to the local pattern.

In addition to comparing the trained condition to the control condition, we performed one-sample t-tests to compare each test condition to 50% chance. In the trained condition, all three test conditions were significantly greater than chance: Old ($t(9) = 4.01, p < 0.01$) and New ($t(9) = 9.46, p < 0.0001$) and First-Order ($t(9) = 2.78, p < 0.050$). The Control condition showed no significant differences from chance in the Old and New Test Items: Old ($t(9) = 0.98, p = 0.37$) and New ($t(9) = 0.90, p = 0.39$), but did show responses significantly below chance for First-Order Test Items ($t(9) = 2.27, p < 0.05$). The below chance performance of the Control participants on First-Order Test Items could represent a bias against first-order harmony if one hears words that begin with a sibilant consonant. If such a bias exists, the results are strengthened because then the trained participants were able to overcome this bias.

Because responses in the trained condition were significantly above chance, the general result that participants generalized to the First-Order Test Items is still valid.

Of the 10 participants in the trained condition, only 2 selected the harmonic option in the First-Order Test Items at a rate of less than 50%. This suggests that while the majority of participants inferred a general harmony pattern, some participants may have inferred a specific second-order harmony pattern.

To further demonstrate that participants in Experiment 2 generalized to the first-order pattern, but participants in Experiment 1 failed to generalize to first-order patterns, we performed a cross-experiment ANOVA comparing responses in the trained conditions of Experiment 1 and Experiment 2. There was no significant effect of Training ($F(1, 18) = 1.43, p = 0.25$), suggesting that the first and second-order long-distance patterns were equally learnable. There was a significant effect of Test Item ($F(2, 36) = 4.69, p < 0.05$), reflecting a difference between Old and New Items and the First/Second-Order Items ($F(1, 18) = 5.39, p < 0.05$). This difference was carried by Experiment 1, as there was a significant interaction ($F(2, 36) = 3.67, p < 0.05$), reflecting a marginally significant difference between First and Second-Order Items between Experiment 1 and 2 ($t(18) = 1.93, p = 0.07$). In addition, the interaction reflects the fact that there was significant difference between Second-Order and New Test Items in Experiment 1 ($t(18) = 2.09, p = 0.05$), but there was no significant difference between First-Order and New Test in Experiment 2 ($t(18) = 0.54, p = 0.60$).

Discussion

Participants in Experiment 2 learned a non-local consonant harmony pattern, and generalized this pattern to first-order instances. This suggests that a second-order long distance harmony pattern implies a first-order long distance pattern, supporting the hypothesis that first-order patterns in harmony are privileged.

GENERAL DISCUSSION

In many ways, consonant harmony appears to be an exception to the generalization that phonological processes are local. Consonant harmony is more likely to be non-local than vowel harmony (i.e., subject to non-adjacent dependencies across a large number of syllables). The present experiments have shown that first-order non-local instances of consonant harmony appear to be privileged over second-order non-local interactions. In an artificial grammar learning situation, learning a second-order non-local consonant harmony pattern implies a first-order consonant harmony pattern, while learning a first-order consonant harmony pattern does not imply a second-order non-local consonant harmony pattern⁴.

Three different proposals for accounting for locality restrictions in vowel harmony were discussed in the Introduction: strict locality, feature-copying, and the tiers approach. In strict locality, second-order and first-order restrictions are represented through coarticulation between the two agreeing segments. Thus, if a language allows second-order long-distance effects in harmony, it means that coarticulation can occur between intervening consonants, which implies that coarticulation will apply between intervening vowels as well. Strict locality keeps a privileged status for first-order representations: second-order implies first-order. The problem with strict locality is that second-order long distance consonant harmony patterns do not appear to adhere to the required coarticulatory demands. As Gafos (1998)

⁴There are cases of non-local patterns in phonology (specifically involving stress) that do not have this implication.

argues, such second-order patterns are best analyzed as feature-copying processes. In basic feature-copying processes, all processes are assumed to be second-order. That is, there are no restrictions as to the distance between the agreeing segments. There are two ways to allow for first-order restrictions on feature-copying consonant harmony. The first is to assume that first-order restrictions follow strict locality, and do not follow under feature-copying analyses. This would account for the results in the present experiment, if it were the case that the stimuli used in first-order pattern were controlled to adhere to strict locality. Because all stimuli were produced in the same manner for both first and second-order locality, it is unlikely that the first-order stimuli truly followed strict locality. This supports the second proposal for accounting for first and second-order restrictions on consonant harmony, which is to include an additional constraint restricting distance. For example, Rose and Walker (2004) propose a 'Proximity' constraint that restricts agreeing segments to vowels, excluding second-order consonant harmony patterns. When Proximity is high-ranked, second-order harmony patterns are excluded. While Proximity allows for both first and second-order constraints in consonant harmony, it does not guarantee that first-order patterns will not imply second-order patterns.

Under the tiers approach, cases of second-order sibilant harmony are captured by placing sibilants on a separate 'sibilant' tier. Adjacent sibilants on the sibilant tier must agree. These sibilants may not be adjacent on the surface, skipping any number of consonants and vowels. A first-order only case of sibilant harmony can be captured if the sibilant agreement rule is stated over the consonant tier: adjacent sibilants on the consonant tier must agree. In this case, if a consonant intervenes between the two sibilants, harmony will not apply. In order for the implication found in the present experiments to apply in the tiers case, learners exposed to first-order patterns must assume that the consonant harmony rule applies at the consonant tier only (not the sibilant tier), implying a first-order pattern, and not a second-order pattern. Learners exposed to second-order harmony must assume that the consonant harmony pattern applies at the sibilant tier and not the consonant tier, allowing for both first and second-order patterns. In this way, there must be a bias against use of tiers for specific feature values (as opposed to consonants only). If a rule can be captured using the consonant tier only (without use of a separate sibilant tier), learners will do so, thereby creating a privilege for first-order patterns.

Both the tiers and the feature-copying approaches to consonant harmony allow a mechanism for representing first-order harmony only. The question raised is to what type of learning algorithm will ensure that the specific, first-order pattern is activated when exposed only to first-order data. Because the standard feature-copying analysis can account for first-order patterns without the Proximity constraint, there is no guarantee that learners will activate the constraint when hearing only first-order patterns. The only way to account for the data in the experiments is if learners in Experiment 1 inferred that the Proximity constraint was active⁵ and therefore did not extend the harmony pattern to second-order non-local instances, while participants in Experiment 2 inferred that the Proximity constraint was not active.

Computational models of learning phonological theories (Boersma & Hayes, 2001; Hayes & Wilson, 2008; Jarosz, 2006; Tesar & Smolensky, 1998) may help to clarify whether different models of consonant harmony show biases towards first-order relationships. For example, if a computational model of Rose and Walker's analysis were trained on the exposure stimuli from Experiment 1, it would be possible to determine what learners infer regarding the Proximity Constraint when exposed only to first-order harmony patterns. Use of such computational models may help to differentiate between models of harmony.

---

⁵In Optimality Theory (Prince & Smolensky, 1993/2004), this constraint would be high-ranked in the first-order case, but low-ranked in the second-order case.

While there exists several learning models that account for second-order non-local phonological processes (Hayes and Wilson, 2008; Heinz, 2010), none of them are designed specifically to account for the implicational universal that a second-order nonlocal pattern implies the existence of a first-order pattern. However, it is possible that such models can ‘explain away’ (Pearl, 1988) the implication. ‘Explaining away’ refers to methods for finding a single cause for an event based on independent evidence. In the case of consonant harmony, a model must decide the source of the harmony pattern: a long-distance pattern that applies to first and second-order patterns, a pattern that is first-order only, or a pattern that is second-order only. If the representations encoded into the model allow only for a first-order and first and second-order harmony patterns (but not second-order patterns only), then when exposed to a second-order pattern, the model will be forced to explain the pattern as a first and second-order pattern only. However, when exposed to first-order information, the model can account for the data as either a first-order pattern or as a first and second-order pattern. The model may choose between these two cases based on additional information, such as a preference for the general, simpler pattern, or a bias towards interpolation over extrapolation. Both heuristics would lead the model to posit a first-order pattern. However, a heuristic towards highly general patterns might lead the model towards the first and second-order pattern. If stochastic mechanisms decide which heuristic is implemented, such that the bias towards highly general patterns has low probability, extrapolation from a first-order pattern a first and second-order pattern would occur with a low frequency, explaining why a small proportion of learners inferred a highly general pattern. Implementing Optimality Theoretic models of learning in a stochastic framework could shed light on how Optimality Theoretic models of harmony, such as Rose and Walker’s (2004) analysis, can account for these inferences.

Work on computational modeling of artificial grammar experiments have been used to help uncover the mechanisms that may be responsible for patterns of learning and generalization of novel linguistic patterns (Perfors, Tenenbaum, & Wonnacott, 2010). In the present experiments, there was a bias for interpolation over extrapolation: generalization to first-order long distance dependencies but second-order. This suggests that a conservative learning algorithm may be most appropriate to capture learning data. However, it is unclear whether all language learning follows a conservative approach. Finley and Badecker (2009a) showed that learners of a vowel harmony pattern were able to extrapolate to novel vowels not heard in training, from low and mid vowels to high vowels. This suggests that some phonological patterns may lead learners to form highly general representations (like distinctive features), allowing for extrapolation.

The present experiment used a suffixing harmony rule. While learners had no trouble learning the suffixing pattern, the majority of languages with sibilant harmony actually have prefixing patterns, or right-to-left spreading (Hansson, 2001). It is a question for future research whether directional constraints play a role in the learning of consonant harmony languages. Given the current results, which showed highly robust learning for a suffixing pattern (as opposed to the more typologically frequent prefixing pattern), it is unlikely that direction of consonant harmony affects learning (at least for adults). Future research will also look at other phonological factors that may affect locality restrictions in consonant harmony, such as the segments involved in harmony, the harmonic feature and syllabic structure.

Further, one dimension of consonant harmony that was largely unexplored in the present paper was similarity. Rose and Walker (2004) show that harmony is most likely to apply if the consonants share several phonological features (and are thus highly similar). This is comparable to the obligatory contour principle () in which similar consonants are disallowed, unless they share the same autosegmental representation (and are thus identical).

Experimental and computational models of learning and representations may help uncover the reasons why similarity plays a role in phonological processing. Previous work on the OCP has used computational approaches to understand how words in languages with the OCP (e.g., Arabic) are restricted (Frisch, Pierrehumbert, & Broe, 2004; Graff & Jeager, in press). Future work may employ computational models of similarity to create a more unified approach to phonological processes that reference featural similarity, such as consonant harmony and the OCP.

A remaining question subject to future research is role of distance in non-local patterns. In the present experiment, all suffixed items were three syllables long and the second-order non-local dependency skipped only a single syllable. It is unclear what the maximum learnable distance is for non-adjacent dependencies. In the present experiment, participants generalized from a non-local pattern to a local pattern, but it is unclear whether participants learned a general non-local pattern that would extend to distances greater than a syllable. Recent research has shown that it is possible to learn consonant harmony patterns with two intervening syllables, suggesting that the consonant harmony pattern used in the present study is not limited to short words (Finley, submitted). Future research will continue to investigate the role of distance in non-local dependencies.

CONCLUSIONS

The present study established the privileged status of local instances of a nonlocal consonant harmony using adult artificial grammar learning experiments. Learners who were exposed to a second-order non-local harmony pattern inferred a first-order version of the harmony pattern, while learners exposed to the first-order version of a consonant harmony pattern failed to infer a second-order non-local pattern. This implication supports the hypothesis that first-order patterns are privileged in harmony. While second-order non-local patterns exist in harmony, the presence of such patterns entails the existence of first-order harmony patterns.

Acknowledgments

This paper would not have been possible without the helpful discussion, advice and suggestions from the following people: Jeffery Heinz, Neil Bardhan, Patricia Reeder, Anna States, Lily Schieber, Carrie Miller, Elissa Newport, helpful and prompt anonymous reviewers, and members of the Aslin-Newport Lab. Portions of this work were presented at the University of Delaware and Rutgers University, Camden. Funding was provided by NIH Grant DC00167. All errors are my own.

References

- Benus S, Gafos AI. Articulatory characteristics of Hungarian ‘transparent’ vowels. *Journal of Phonetics*. 2007; 35:271–300. [PubMed: 18389086]
- Boersma P, Hayes B. Empirical tests of the gradual learning algorithm. *Linguistic Inquiry*. 2001; 32:45–86.
- Chomsky, N. *Lectures on government and binding*. Dordrecht: Foris; 1981.
- Cohen JD, MacWhinney B, Flatt M, Provost J. *PsyScope: A new graphic interactive environment for designing psychology experiments*. *Behavioral Research Methods, Instruments and Computers*. 1993; 25:257–271.
- Culicover, P.; Wilkins, W. *Locality in linguistic theory*. New York: Academic Press; 1984.
- Finley, S. Testing the limits of long-distance learning: Learning beyond the three-segment window. submitted
- Finley S, Badecker W. Analytic biases for vowel harmony languages. *WCCFL*. 2008; 27:168–176.
- Finley S, Badecker W. Artificial grammar learning, and feature-based generalization. *Journal of Memory and Language*. 2009a; 61:423–437.

- Finley, S.; Badecker, W. Right-to-left biases for vowel harmony: Evidence from artificial grammar. In: Shardl, A.; Walkow, M.; Abdurrahman, M., editors. Proceedings of the 38th North East Linguistic Society Annual Meeting. Vol. 1. 2009b. p. 269-282.
- Finley S, Badecker W. Towards a substantively biased theory of learning. BLS. :33. in press.
- Frisch SA, Pierrehumbert J, Broe MB. Similarity avoidance and the OCP. *Natural Language and Linguistic Theory*. 2004; 22:179–228.
- Gafos AI. Eliminating long-distance consonantal spreading. *Natural Language & Linguistic Theory*. 1998; 16(2):223–278.
- Gafos, AI. The articulatory basis of locality in phonology. New York: Garland; 1999.
- Goldsmith, J. Autosegmental phonology. Garland Press; 1979.
- Graff, P.; Jeager, TJ. Locality and feature specificity in OCP effects: Evidence from Aymara, Dutch, and Javanese. Proceedings of the 45th meeting of the Chicago Linguistics Society; Chicago, IL: CLS; in press
- Hansson, GO. PhD dissertation. University of California; Berkeley, Berkeley: 2001. Theoretical and typological issues in consonant harmony.
- Hayes B, Wilson C. A Maximum Entropy Model of Phonotactics and Phonotactic Learning. *Linguistic Inquiry*. 2008; 39:379–440.
- Heinz J. Learning long-distance phonotactics. *Linguistic Inquiry*. 2010; 41(4):623–661.
- Hyman, LH. Phonologization. In: Julliland, A., editor. *Linguistic studies offered to Joseph Greenberg*. Vol. 2. Saratoga, CA: Anma Libri; 1976. p. 407-418.
- Jarosz, G. Richness of the base and probabilistic unsupervised learning in Optimality Theory. Proceedings of the eighth meeting of the ACL special interest group in computational phonology; 2006.
- Keating P. Underspecification in phonetics. *Phonology*. 1988; 5(2):275–292.
- McDonough, J. The Navajo sound system. Dordrecht: Kluwer Academic Publishers; 1990.
- Nevins A. On formal universals in phonology. *Behavioral and Brain Sciences*. 2009; 32:461–462.
- Nevins, A. Locality in vowel harmony. Cambridge, MA: MIT Press; 2010.
- NiChiosain; Padgett, J. Markedness, segment realization and locality in spreading. In: Lombardi, L., editor. *Segmental phonology in Optimality Theory*. Cambridge: Cambridge University Press; 2001. p. 118-156.
- Ohala, JJ. Towards a universal, phonetically-based theory of vowel harmony. 3rd International Conference on Spoken Language Processing; 1994. p. 491-494.
- Öhman SEG. Coarticulation in VCV utterances: Spectrographic measurements. *Journal of the Acoustical Society of America*. 1966; 39:151–168. [PubMed: 5904529]
- Pearl, J. Probabilistic reasoning in intelligent systems: Networks of plausible inference. San Mateo, Ca: Morgan Kaufman; 1988.
- Perfors A, Tenenbaum JB, Wonnacott E. Variability, negative evidence, and the acquisition of verb argument constructions. *Journal of Child Language*. 2010; 37:607–642. [PubMed: 20367896]
- Prince, A.; Smolensky, P. *Optimality theory: Constraint interaction in generative grammar*. Cambridge: Blackwell; 1993/2004.
- Pycha A, Nowak P, Shin E, Shosted R. Phonological rule-learning and its implications for a theory of vowel harmony. *WCCFL*. 2003; 22:101–113.
- Ringen CO, Vago RM. Hungarian vowel harmony in Optimality Theory. *Phonology*. 1998; 15:393–416.
- Rose S, Walker R. A typology of consonant agreement as correspondence. *Language*. 2004; 3:475–531.
- Sapir, E.; Hojier, H. The phonology and morphology of the Navajo language. Vol. 50. University of California Publications; 1967.
- Tesar B, Smolensky P. Learnability in Optimality Theory. *Linguistic Inquiry*. 1998; 29:229–268.
- Vilain, A.; Abry, C.; Badin, P. Coproduction strategies in French VCVS: Confrontin Ohman's model with adult and developmental articulatory data. *Proc. of the 5th Seminar on Speech Production: Models and Data; Kloster Seeon*. 2000. p. 81-84.

- Wilson C. Experimental investigations of phonological naturalness. WCCFL. 2003; 22:101–114.
- Wilson C. Learning phonology with substantive bias: An experimental and computational study of velar palatalization. Cognitive Science. 2006; 30:945–982.

Appendix 1: Experiment 1 Stimuli

Training Items: piʃa-piʃaʃu, poʃe-poʃeʃu, peʃu-peʃuʃu, tiʃo-tiʃoʃu, tuʃa-tuʃaʃu, toʃi-toʃiʃu
koʃa-koʃaʃu, kuʃo-kuʃoʃu, kiʃe-kiʃeʃu, biʃo-biʃoʃu, boʃa-boʃaʃu. beʃa-beʃaʃu, daso-
dasosu, dusi-dusisu, diso-disosu, gosa-gosasu, gesu-gesusu, gaso-gasosu, muso-musosu,
mase-masesu, mesu-mesusu, nasu-nasusu, nesi-nesisu, neso-nesosu

Old	beʃaʃu, boʃaʃu, disosu, dusi, gesu, kuʃoʃu, masesu, mesu, niʃaʃu, piʃaʃu, poʃeʃu, tuʃaʃu
New	nesosu, deʃeshu, geʃiʃu, basosu, desasu, gusasusu, keshiʃu, kishuʃu, miʃuʃu, paʃisu, pisesu
Second-Order	ʃanoʃu, ʃeboʃu, ʃetaʃu, ʃikaʃu, ʃikuʃu, ʃokiʃu, ʃomiʃu, ʃupaʃu, ʃapiʃu, ʃegeʃu, ʃideʃu, ʃuboʃu

Test Items (Harmonic Items):

Appendix 2: Experiment 2 Stimuli

Training Items: suge-sugesu, sone-sonesu, ʃupe-ʃupe, ʃoku-ʃokuʃu, suna-sunasu, soga-
sogasu, ʃako-ʃakoʃu, sume-sumesu, ʃobi-ʃobiʃu, ʃito-ʃitoʃu, sodi-sodisu, somu-somusu,
ʃepo-ʃepoʃu, ʃabo-ʃaboʃu, ʃapi-ʃapiʃu, sago-sagosu, sidu-sidusu, sodu-sodusu. sema-
semasu, ʃoti-ʃotiʃu, sine-sinesu, ʃabe-ʃabeʃu, ʃeki-ʃekiʃu, ʃatu-ʃatuʃu

Test Items (Harmonic Items):

Old	semasu, ʃabeʃu, ʃakoʃu, ʃekiʃu, ʃitoʃu, ʃokuʃu, ʃupeʃu, sidusu sodisu, sogusu, sugesu, sumesu
New	ʃanoʃu, ʃeboʃu, ʃetaʃu, ʃikaʃu, ʃikuʃu, ʃokiʃu, ʃomiʃu, ʃupaʃu, ʃapiʃu, ʃegeʃu, ʃideʃu, ʃuboʃu
First-Order	beʃaʃu boʃaʃu, deʃeʃu, geʃiʃu, keʃuʃu, kuʃoʃu, miʃuʃu, niʃaʃu, paʃiʃu, piʃaʃu, poʃeʃu, tuʃaʃu

Figure 1.
Experiment 1 Results: Means and Standard Errors

Figure 2.
Experiment 2 Results: Means and Standard Errors

Table 1

Representation of Locality

Theory	Representation										Locality Restrictions	
Strict Locality Consonant-Vowel Tiers	s				a _s				t _s	i _s	s	Dependent on Degree of Coarticulation First-Order Non- Locality Only
	s				t					C-Tier V-Tier		
Sibilant-Consonant-Vowel Tiers	s											None
	s				t			s		Sib-Tier C-Tier V-Tier		
Feature Copying		a										None
	s ₁				a				t	i	s ₁	

Table 2

Examples of Training Stimuli

Experiment 1: First-Order Non-Locality	Experiment 2: Second-Order Non-Locality
beʃa, beʃaʃu	ʃabe, ʃabeʃu
tuʃa, tuʃaʃu	ʃito, ʃitoʃu
mesu, mesusu	sogu, sogusu
dusi, dusisu	sema, semasu

Table 3

Examples of Test Stimuli

Experiment 1		
Old	New	Second-Order
beʃaʃu vs. *beʃasu	desasu vs. *desaʃu	ʃanoʃu vs. *ʃanosu
tuʃaʃu vs. *tuʃʃasu	nesosu vs. *neoʃu	ʃeteʃu vs. *ʃetasu
dusisu vs. *dusiʃu	geshiʃu vs. *geʃʃisu	ʃuboʃu vs. *ʃubosu
Experiment 2		
Old	New	First-Order
ʃabeʃu vs. *ʃabesu	ʃanoʃu vs. *ʃanosu	beʃaʃu vs. *beʃasu
ʃitoʃu vs. *ʃitosu	ʃeteʃu vs. *ʃetasu	tuʃaʃu vs. *tuʃʃasu
semasu vs. *semaʃu	ʃuboʃu vs. *ʃubosu	dusisu vs. *dusiʃu