Analytic Biases for Vowel Harmony Languages

Sara Finley and William Badecker
Johns Hopkins University

1. Introduction[1]

One of the major goals of generative linguistics is to produce a theory of language that can generate all possible languages without predicting languages that are unnatural in that they lie outside the scope of the human capacity to learn. In order to achieve this goal, the nature of language must be understood with regard to the distinction between linguistic patterns that are outside the range of possible grammars and patterns that are accidental gaps, and simply have not been documented. This difference becomes crucial when evaluating theories of generative linguistics that predict linguistic patterns that have never been observed in the set of natural languages. Specifically, grammars generated within Optimality Theory (OT) (Prince and Smolensky 1993/2004) produce a factorial typology of all linguistic patterns predicted by the given grammatical theory. All optimality theoretic grammars are predicted to be within the cognitive capacity of language users.

The substantively biased theory of learning (Finley and Badecker, 2008; Wilson 2006) provides a promising means for understanding the relationship between grammars generated by linguistic theories and the characterization of patterns observed cross-linguistically. This theory of language learning hypothesizes that learning biases shape the distribution of linguistic patterns across the world's languages. The easiest patterns to learn are consequently the most common cross-linguistically. Patterns that are phonetically grounded and/or formally concise are the easiest patterns to learn, and therefore the most likely to appear cross-linguistically, while patterns that lack these properties are avoided by the learner and are therefore cross-linguistically rare. In phonology, these learning biases are grounded in both phonetic naturalness as well as phonological naturalness. For example, learners are biased to form grammars that maximize perceptual salience and articulatory ease, but they will also be parsimonious in terms of the formalization of the grammar. These formal restrictions are characterized in terms of grammatical constructs such as natural classes and formal implementation (e.g., number of rules or constraints required to characterize the grammar); they may also include non-linguistic factors that influence language processing such as working memory and attention.

While the substantively biased theory of learning offers a means for explaining the relationship between frequently occurring and unattested linguistic patterns, there is little concrete evidence to support the notion that learning biases shape the cross-linguistic distribution of patterns in the world's languages. Specifically, traditional methods for understanding the nature of linguistic typologies are limited to exploring attested patterns, and typically focus on frequently occurring or natural patterns. Because it is impossible to study how learners in a natural setting will cope with an unattested pattern, it is unclear why learners ultimately avoid particular unattested patterns. It may be that with exposure to the proper learning data, learners may be equally accommodating toward outwardly unnatural, unattested patterns as frequently occurring, natural patterns.

Because traditional methods cannot address the ways in which learners interpret unattested patterns, this paper employs the artificial grammar learning paradigm (Finley and Badecker, in press; Reber, 1967; Wilson 2006) in order to address the question of how learners deal with data that is ambiguous between naturally occurring patterns and unattested patterns. In the artificial grammar learning paradigm, the experimenter can control the data that the learner is exposed to, making it possible to investigate the nature of learning biases towards natural versus unnatural patterns as well as

[1] We are grateful for the following people who have provided helpful comments, advice and time: Iris Berent, Colin Wilson, Paul Smolensky, Ariel Goldberg, Elliot Moreton, Bruce Hayes, Deepti Ramadoss, Rebecca Morley, Manny Vindiola, members of the JHU Psycholinguistics Lab, members of the JHU IGERT lab, and our audience at WCCFL 27. All errors are our own.

Proceedings of the 27th West Coast Conference on Formal Linguistics, ed. Natasha Abner and Jason Bishop, 168-176. Somerville, MA: Cascadilla Proceedings Project.

attested versus unattested patterns. According to the substantively biased learning hypothesis, a learner who is presented with language data that is ambiguous between a naturally occurring pattern and an unnatural or unattested pattern will infer the naturally occurring pattern. In the artificial grammar learning setting, it is possible to directly manipulate the learning data such that it contains precisely this ambiguity. If learners conform to substantive biases, when confronted with data that is ambiguous between an attested pattern and unattested pattern, they will postulate an attested pattern.

This paper tests the assumptions of the substantively biased learning hypothesis in terms of vowel harmony, a phonological pattern in which all vowels in a lexical item share the same feature value (e.g., if the first vowel in a lexical item is front ([i, e]) all other vowels in the word must also be front ([i, e])). Vowel harmony is an ideal phonological pattern to test the substantively biased learning hypothesis for several reasons. First, vowel harmony is an extremely common phonological pattern that occurs in a wide range of language families (e.g., Bantu, Nilotic, Romance, Uralic) (Clements, 1976; Kiparsky 1981). Second, vowel harmony is widely studied and the typology of vowel harmony is relatively well agreed upon among linguists, making it possible to assess which vowel harmony patterns are natural and which are unattested (Nevins, 2004). Finally, vowel harmony is ideal for an artificial grammar learning experiment because English does not have vowel harmony, and is relatively exotic to the typical native speaker of American English. Vowel harmony can therefore be taught to our participants without worry of prior knowledge or obscurity of the phonological pattern we are testing.

Vowel harmony is a spreading process according to which a single vowel will serve as the harmony trigger- and will thereby determine the feature value of all other vowels in the word. Apart from morphological influences on vowel harmony, there are two main factors that determine how the trigger for harmony is determined: directionality and dominance (Bakovic 2000; Nevins, 2004)[2]. In directional harmony systems, either the leftmost vowel spreads its features rightward (e.g., /+ – –/ → [+ + +]; /– + +/ → [– – –]) or the rightmost vowel spreads its feature value leftwards (e.g., /+ + –/ → [– – –]; /– – +/ → [+ + +]). In dominant systems, the harmony trigger is determined by the presence of a vowel with a particular feature value (e.g., [+F]). If such a vowel is present, the feature will spread bi-directionally (e.g., /– + –/ → [+ + +]).

However, there are no languages in which the harmony trigger is determined by the relative number of vowels in the input with a particular feature value. This type of unattested pattern, termed 'majority rules' (Lombardi 1998), hypothetically occurs when the trigger for harmony is determined by the feature value in the input that will allow for all vowels to share the same feature value, but with the fewest possible changes to the input. For example, if the input contains two [+F] vowels and one [–F] vowel, the output will have three [+F] vowels (/+ – +/ → [+ + +]), but if the input contains two [–F] vowels and one [+F] vowel, the output will be [–F] (e.g., /– – +/ → [– – –]), regardless of the direction of harmonic spreading.

While such patterns are unattested, they are very easily generated in grammatical systems like Optimality Theory (Bakovic 1999, 2000; Lombardi 1998). With a constraint that only requires output vowels to share the same feature value (e.g., AGREE[F]), it is possible, with a markedness, faithfulness interaction to predict a language in which all outputs are harmonic, but the harmonic trigger is determined by the majority feature value, because that majority feature value will produce the fewest violations of faithfulness. This is illustrated in the hypothetical tableau in (1).

(1) AGREE >> ID

/bɪnigi/	AGREE[ATR]	ID[ATR]	*[+ATR]	*[–ATR]
[bɪnigi]	*!		**	*
[bɪnɪgɪ]		**!		***
☞ [binigi]		*	***	

[2] Another type of harmony pattern, not discussed, is stress-triggered harmony in which the feature value of the stressed vowel spreads throughout the lexical item. This can be considered a form of dominance because spreading is bi-directional.

The uniformly [+ATR] candidate [binigi] is the winning output of the optimization in (1) because it has the fewest ID[ATR] violations while still satisfying the higher-ranked AGREE constraint: there are two [+ATR] vowels and only one [–ATR] vowel in the input. Had the input contained two [–ATR] vowels and one [+ATR] vowel (e.g., /bɪnɪgi/), the optimal output would have been the [–ATR] candidate [bɪnɪgɪ]. The markedness-faithfulness interaction of OT, along with a harmony constraint that does not require dominance or directionality, can produce harmony systems in which any harmonic output will surface with the least number of changes to the input. If the harmony inducing constraint only requires agreement, then the harmonic item with the fewest faithfulness violations will emerge, producing an unattested majority rules grammar.

The question, then, is whether the fact that majority rules grammars are easy to produce should be a problem for the analyst. While Lombardi (1998) and Bakovic (1999, 2000) have found solutions to avoiding majority rules grammars in analyses that make use of non-directional AGREE-type harmony constraints, there is still the worry that the ease of generation of such grammars implies that majority rules grammars should be available to the cognitive learning apparatus. It is also possible that majority rules grammars may be entirely plausible from the standpoint of the grammatical system and the fact that there are no majority rules languages is due to an accidental gap in grammars that can produce such majority rules languages.

One way to distinguish between a principled restriction on the nature of vowel harmony languages and an accidental gap account is through testing for learning biases for both majority rules languages and attested harmony languages, specifically directionally-induced harmony languages. If learners are biased against majority rules languages in favor of attested directional harmony system, then it suggests that the non-existence of majority rules languages is a real restriction on grammars. Because it is impossible to test learning biases for unattested languages in a naturalistic setting, as there are no naturalistic settings where a majority rules grammar might be present, the artificial grammar learning paradigm is the best method for addressing this question.

2. The Experiment

The present experiment tests whether learners infer a majority rules pattern when given training data that are ambiguous between a majority rules harmony pattern and a directional harmony pattern. If learners are biased towards directional patterns and against majority rules patterns, they should infer a directional pattern when given data that are ambiguous between majority rules and fixed directionality. In this experiment, participants are exposed either to a left-to-right harmony pattern or a right-to-left harmony pattern in which the majority of the vowels in the input spread their values for the features [Back] and [Round]. If participants learn a majority rules pattern, they will reverse the direction of spreading when the majority of vowels reverse; but if participants learn a directional pattern, they will select the minority pattern as long as the direction of spreading remains the same.

There are several reasons that this experiment tests for responses to ambiguous data as opposed to comparing learning rates for attested versus unattested languages. First, pilot data suggested that it possible to learn the unattested majority rules pattern, particularly if the training and test stimuli was presented in a way that would utilize particular explicit memory and attention strategies. Second, the substantively biased learning hypothesis states that learners are specifically drawn towards natural linguistic patterns and that given a choice between a natural and unnatural pattern, the learner will hypothesize the natural pattern. Therefore, a true test of this hypothesis will expose learners to such ambiguous stimuli. Finally, because learners are able to acquire unnatural patterns in natural settings (Anderson, 1981), a result in which learners can learn both a natural and an unnatural pattern does not rule out the possibility that learners have a bias towards natural grammars.

2.1. Participants

All participants were adult native American English speakers who had no prior knowledge of a vowel harmony language, and had not participated in any previous vowel harmony learning experiments. Thirty-six Johns Hopkins undergraduate students participated for extra course credit.

Participants were randomly assigned to one of three training conditions: a Control Condition containing disharmonic input forms, a Right-to-Left Condition and a Left-to-Right Condition. Final analyses included 11 participants in each of three conditions. All participants were screened based on a perceptual (AXB) task, described below. Those participants scoring less than 75 percent on this task were removed from the study. This occurred for 1 participant. The data for 2 participants were not used because they had been run in previous vowel harmony learning experiments.

2.2. Design

Participants in the Right-to-Left and Left-to-Right conditions were exposed to a vowel harmony language in which all vowels agreed in the features round and back, and were drawn from the set of English vowels that included [i, e] (front, unround) and [u, o], (back, round). All vowel alternations kept vowel height distinctions intact such that [i] alternated with [u] and [e] alternated with [o].

While previous experiments exposed participants to a harmony alternation in the form of a stem-suffix alternation (Finley & Badecker 2008; in press), this experiment used morphophonological alternations to the display the harmony pattern. In this paradigm, participants were presented with a single disharmonic 3-syllable word followed by its harmonic counterpart (e.g., [pigebo] → [pigebe]). The disharmonic input was presented to the participants in order to make it explicit as to which vowels were changing from the input to the output, making it possible to encode the fact that all items followed a majority rules pattern in which the minority feature value was always the vowel that changed. Participants in the Right-to-Left Condition heard items that were ambiguous between majority rules and right-to-left spreading while participants in the Left-to-Right Condition heard items that were ambiguous between majority rules and left-to-right spreading. Participants in the Right-to-Left and Left-to-Right conditions each heard 24 sets of disharmonic-harmonic alternations, and were trained on either right-to-left harmony (e.g., [pidego] → [pidege]) or left-to-right harmony (e.g., [pidogo] → [pudogo]), respectively.

Because the experiments are testing the nature of learning in an artificial grammar setting, it is necessary to ensure that the effects that we find are actually due to learning, rather than a prior bias, or bias in selecting particular items (Redington & Chater 1996). We created control conditions to ensure that all effects were due to specific learning strategies. One option for a Control condition is to give these participants no training data at all, and simply give test items to these participants. However, we do not consider this to be an adequate control for the present experiment. Insofar as these control participants would have no training on which to base their responses, participants would require very different test instructions in order to have a sense for how to perform the task. This variation in instructions may influence the strategy for responding, as well as the results. The fact that they would have no basis in learning for their responses could lead participants to use a strategy that is entirely different from any strategy that participants in the critical training conditions might employ. Our solution to these concerns was to give the participants in the Control condition the underlying forms (disharmonic items). This would provide them a sense of how the language sounds, but no basis for how the harmonic items should sound in the alternation.

All stimuli were recorded in a sound proof booth at 22,000kHz by a male speaker of American English with basic phonetic training (had completed a graduate-level phonetics course). While the speaker had no knowledge of the specifics of the experimental design, he was aware that the items would be used in an artificial language learning task. All stimuli were phonetically transcribed, and presented to the speaker in written format. The speaker was instructed to produce all vowels as clearly and accurately as possible, even in unstressed positions. Stress was produced on the first syllable.

All sound editing was done using Praat (Boersma & Weenink, 2005). All stimuli contained the same consonant inventory: [p, b, t, d, k, g, m, n]. The vowel inventory for all conditions consisted of [i, u, e, o]. The training stimuli were counterbalanced to contain all possible combinations of vowel sounds. Consonants were also counterbalanced so that all consonants appeared equally often in each position. Stimuli were created semi-randomly with the condition that any word too closely resembling an English word was intentionally avoided (the final profile of the stimuli was counterbalanced to appropriately contain equal numbers of consonant pairs and a consistent number of vowel pairs). Examples of training stimuli appear in the table in (2).

(2) Example Training Stimuli

Right-to-Left Training Condition	Left-to-Right Training Condition
[bedutu, bodutu]	[boduti, bodutu]
[dogibe, degibe]	[degibo, degibe]
[gebupu, gopubu]	[gopubi, gopubu]
[mogite, megite]	[megito, megite]

All harmony alternations were presented sequentially such that the harmonic alternation occurred immediately following the disharmonic input base form. For example, participants in the Left-to-Right Condition heard a disharmonic form immediately followed by its harmonic counterpart, with the final vowel alternating (e.g., [boduti, bodutu]).

Test items were presented in terms of two pairs of items, in which the first item in each pair was a disharmonic and the second item was its harmonic counterpart. Participants were asked to choose between two alternations of the same disharmonic input, one with spreading from right-to-left (e.g., [pumite, pumuto]), and the other with spreading from left-to-right (e.g., [pumite, pimite]). The order of paired test items was counterbalanced. Both pairs were heard sequentially with one immediately followed by the other (e.g., [pumite, pimite], [pumite, pumuto]) and participants were asked to select which pair was most likely to be from the language they heard in training.

Critical test items were reversed such that spreading the majority feature value requires spreading in the opposite direction from what was contained in the training set, one of which contained all round vowels, and the other contained all unround vowels. If learners inferred a directional pattern, then they should accept multiple items undergoing harmony from the input to the output. If learners inferred a majority rules pattern, they should reverse the direction of spreading. Test items included 12 Old Items, 12 New Items and 12 New Direction Items. The position of the vowels with the majority back/round feature values in the Old and New Items corresponds to the trigger location/direction of spreading that the participant was trained on, but the New Direction Items reflected a reversal of the direction that the participants were trained on. Representative examples of the test items are provided in (3).

(3) Examples of Test Items

Right-to-Left Training **Old Items**	**Left-to-Right Training** **New Direction Items**
[pumite, pimite]	[pumite, pumuto]
[kogibe, kegibe]	[kogibe, kogubo]
[nepoko, nopoko]	[nepoko, nepeke]
New Items	**New Items**
[nupiki, nipiki]	[nipiku, nipiki]
[pedobo, podobo]	[podobe, podobo]
[biteko, biteke]	[biteko, biteke]
New Direction Items	**Old Items**
[kukope, kikepe]	[kukope, kukopo]
[miketo, mukoto]	[miketo, mikete]
[nebegu, nobogu]	[nebegu, nebegi]

Participants in all conditions were given virtually identical test items. The major difference was that Old Items in the Left-to-Right Condition were New Direction items for the Left-to-Right Condition.

2.3. Procedure

The experiment consisted of a training phase followed immediately by a forced-choice grammaticality test. All phases of the experiment were presented using PsyScopeX (Cohen, MacWhinney, Flatt, & Provost, 1993). All participants were given written and verbal instructions. They were told that they would be listening to a language they had never heard before, and that they would later be asked about the language, but they need not try to memorize any forms they heard. No information about vowel harmony or about the morphology of the language was given. No semantics accompanied the sound pairs. Participants heard all stimuli over headphones. Participants were presented with the underlying disharmonic form before the harmonic form. There were 24 items for each participant, each repeated 5 times in a random order.

Training was followed by a forced-choice test phase in which participants heard the disharmonic input followed by a choice of harmonic alternations: all round or all unround. If the first alternation belonged to the language, participants were instructed to push the 'a' key on the keyboard; if the second alternation belonged to the language, participants were instructed to press the 'l' key on the keyboard. Participants were told to respond as quickly and accurately as possible, and to make their responses after hearing both alternations.

The experiment finished with an AXB perception task. All participants were screened based on this perceptual task to ensure that they could discriminate between round and unround vowels (of the same height, high and mid counterparts). In this simple AXB perception task, participants were presented with three syllables and were asked to judge whether the vowel in the second syllable was identical to the vowel in the first or the third syllable. For example, if participants heard [pe], [ke] and [bo], the correct response would be [pe]. All items were monosyllabic. The task was designed as a way to screen participants for the ability to perceive round-unround contrasts in English. Those participants scoring less than 75 percent on this task were removed from the study, with the logic that we assume that all participants should have general competence for perceiving English vowels in order for their learning data to be meaningful.

No explicit feedback was given to participants following completion of the training phase of the experiment. Although any artificial language task contains some level of arbitrariness, leaving out explicit feedback makes the learning as close to a natural setting as possible and helps to minimize the use of explicit learning strategies. Participants were given a debriefing statement upon completion of the experiment (which took approximately 15 minutes).

2.4. Results

The proportion of 'majority rules' responses was recorded for each participant. If participants learned a majority rules pattern, this proportion should remain high for all test conditions. However, if participants learned a directional pattern, then the proportion of majority rules responses should be high for Old and New test items, but low for New Direction Items. A summary of the results across training and response conditions is provided below in the figure in (4).

To ensure that there was an overall effect of training for each of the conditions, a 2 (Training) x 3 (Test Condition) mixed-design ANOVA was performed, comparing the Left-to-Right and the Right-to-Left conditions with the Control condition. There was a significant effect of Training for both the Left-to-Right condition (mean = 0.57 training vs. 0.42 control, $CI = \pm 0.06$; $F(1, 20) = 11.83$, $p < 0.01$) as well as the Right-to-Left condition (mean = 0.65 training vs. 0.42 control, $CI = \pm 0.06$; $F(1, 20) = 36.80$, $p < 0.001$). Both conditions learned some form of the harmony alternation.

To test whether participants inferred a directional rule versus a majority rules pattern, we performed contrasts comparing the New Direction test condition to the Old and New items respectively. Responses in the Left-to-Right condition also showed a significant difference between the New Direction and both the Old (mean = 0.78 vs. 0.23, $CI = \pm 0.30$; $F(1, 10) = 20.82$, $p < 0.01$) and New (mean = 0.70 vs. 0.23, $CI = \pm 0.27$ $F(1, 10) = 11.86$, $p < 0.01$) test conditions. In the Right-to-Left Condition, there was a significant difference between the New Direction and both the Old test conditions (mean = 0.86 vs. 0.25, $CI = \pm 0.26$; $F(1, 10) = 34.22$, $p < 0.001$) and New test conditions (mean = 0.83 vs. 0.25, $CI = \pm 0.24$; $F(1, 10) = 24.74$; $p < 0.01$). The fact that participants chose the 'majority rules' items significantly more often in the Old and New test items suggests that learners infer a directional pattern rather than a majority rules pattern, which indicates a bias *against* majority rules.

(4) Experiment Results

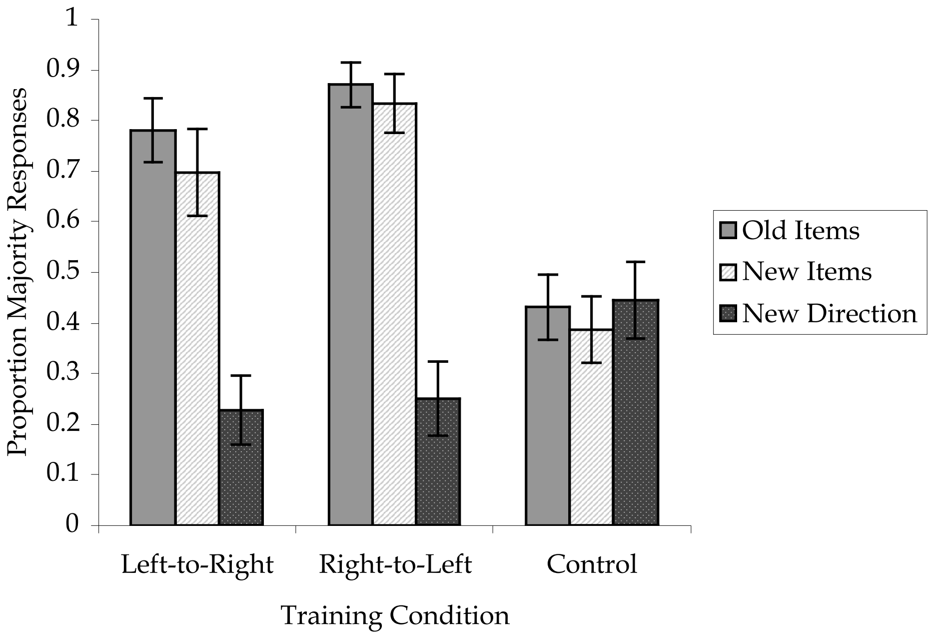

2.5. Discussion

The results of this experiment suggest that learners are biased toward a directional harmony pattern over a majority rules pattern. This was true for both right-to-left as well as left-to-right training conditions, and there was no discernable pattern in the control condition, suggesting that there is a learning bias towards directionally based harmony over majority rules harmony systems.

3. General Discussion

The results of the present experiment suggests that learners exhibit a bias that favors the naturally occurring directionally based spreading pattern over the unattested majority rules spreading pattern. When presented with training data that was ambiguous between a majority rules pattern and a directional spreading pattern, learners overwhelmingly inferred the directional pattern. The fact that learners are biased towards directional patterns over majority rules harmony patterns gives us some insight into why majority rules patterns do not exist in natural language. If learners are not biased to infer majority rules from their language data, then it is unlikely that such a pattern would emerge in a natural setting. Rather, the pattern to emerge would be a directional pattern, which is what learners are biased to infer. This is particularly likely if learning data is ambiguous between the majority rules and some other pattern. Given that morphological factors (e.g., suffix allomorphs) are likely to increase this ambiguity, it is unlikely that a learner will receive the type of unambiguous evidence necessary to infer an unnatural majority rules pattern.

If majority rules grammars were to exist in natural language, it would have to emerge though historical change in languages that originally conform to some other (e.g., directional) generalization governing harmony. For this to occur, the learner would have to infer a majority rules grammar out of ambiguous data, and reanalyze any disconfirming data as majority rules. Because the learner must postulate a grammar that can deal with any number of vowels in the input, majority cases will only comprise a fraction of the possible cases the grammar must account for. An emergent majority rules grammar must therefore posit a default strategy in addition to the majority rules grammar in order to deal with inputs in which there are equal numbers of opposing feature values (e.g., two [+F] vowels and two [–F] vowels). Such a default strategy (e.g., spread leftwards) can easily apply to 'uneven' cases as well as 'tie' cases. This may add an unnecessary level of complication to the grammar as well as requiring additional attentional and memory resources for the learner. By hypothesis, the learner needs only to posit the default pattern for 'tie' cases in order to explain all of the data, as positing a majority rules grammar provides no advantage for explaining the data. Because there is no learning advantage to positing a majority rules grammar, it is unlikely that a learner will postulate an emergent majority rules grammar.

Another potential factor influencing the bias towards directionality over majority rules may be the possibility that learners pay more attention to direction of spreading or location of the undergoer rather than the number of harmony triggers. For example, in our training data, the vowel to undergo harmony was always at either the left or right edge of the word. If learners are able to focus on this aspect of the pattern, they will be less likely to infer a majority rules pattern because their attention will be pulled towards the undergoers of harmony rather than the harmony triggers. In natural language, directional harmony is also structured in this way; in right-to-left harmony the first vowel always alternates; in left-to-right harmony, suffix vowels always alternate at the end of the word. If learners attend to the location of alternations (edge of word, suffix, prefix, etc.), then they will be less likely to remember the number of vowels that triggered harmony, and therefore should be unlikely to infer a majority rules pattern.

The fact that the present experiment demonstrated a bias towards a directional pattern over a majority rules pattern supports a theory of grammar that avoids predicting majority rules grammars. Not only are such languages un-attested, but learners have no bias to infer such patterns. There is a question, however, as to whether majority rules grammars really have to be excluded by the formal grammatical framework[3]. If non-linguistic factors drive the bias against majority rules, then it remains possible that the generative system could produce such majority rules grammars but that language learners will be very unlikely to infer such a grammar as a consequence of non-linguistic constraints on memory and attention. In the present experiments, a directional harmony pattern makes use of the attentional salience of the beginnings and ends of words (Beckman 1998, Brown & McNeill 1966, Christophe, Peperkamp, Pallier, Block & Mehler 2004, Cole 1973, Marslen-Wilson 1975). For example, in left-to-right spreading, the first vowel of the word is always the trigger, and the final vowel of the word is always a target. If learners are biased to infer patterns based on what occurs at the beginnings and ends of words, it is likely that they will infer a directional pattern over a majority rules pattern (which is completely unreliable in terms of salience of the beginnings and ends of the words).

One way that one might be able to determine whether the directionality preference is due to non-linguistic factors against majority rules is to replicate the present experiments using non-linguistic stimuli. If learners of non-linguistic pattern follow the same constraints on majority rules, then it is likely that the bias against majority rules found in these experiments is due to non-linguistic factors. On the other hand, if no bias is found in non-linguistic stimuli, this would suggest that there is something specific to the way the humans perceive and represent language about the nature of harmony which biases learners towards directional spreading.

4. Conclusion

The present experiment supports the substantively biased learning hypothesis that learners are biased towards natural patterns, and that given ambiguous data, learners will infer a natural grammar. The naturalness of directional patterns over majority rules may be based in the fact that learners are

[3] For a formal proposal on how grammars may restrict unattested vowel harmony patterns, see Finley (2008).

sensitive to the location of the undergoer rather than the quantity of harmony triggers. Future research will explore the nature of sensitivity to location and quantity of the harmony undergoer and trigger in learning by specifically controlling for these factors in the training.

References

Anderson, S. (1981). Why phonology isn't "natural". Linguistic Inquiry, 12, 493-547.

Bakovic, E. (1999). Assimilation to the unmarked. Penn Working Papers in Linguistics, 6(1).

Bakovic, E. (2000). Harmony, dominance and control. Unpublished doctoral dissertation, Rutgers University.

Beckman, J. M. (1998). Positional faithfulness. UMass Amherst.

Boersma, P., & Weenink. (2005). Praat: Doing phonetics by computer.

Brown, R., & McNeill, D. (1966). The "tip-of-the-tongue" phenomenon. Journal of Verbal Learning and Verbal Behavior, 5, 325-337.

Clements, G. N. (1976). The autosegmental treatment of vowel harmony. Phonologica, 111-119.

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.

Cole, R. A. (1973). Listening for mispronunciations: a measure of what we hear during speech. Perception and Psychophysics, 11, 153-156.

Christophe, A., Peperkamp, S., Pallier, C., Block, E., & Mehler, J. (2004). Phonological phrase boundaries constrain lexical access I. Adult data. Journal of Memory and Language, 51, 523-547.

Finley, S. (2008). Formal and cognitive approaches to vowel harmony. PhD Dissertation. Department of Cognitive Science, Johns Hopkins University.

Finley, S., & Badecker, W. (2008). Towards a substantively biased theory of learning. BLS, 33.

Finley, S., & Badecker, W. (in press). Right-to-left biases for vowel harmony: Evidence from artificial grammar. NELS, 38.

Kiparsky, P. (1981). Vowel harmony. Unpublished manuscript, MIT.

Lombardi, L. (1999). Positional faithfulness and voicing assimilation in optimality theory. NLLT, 17, 267-302.

Marslen-Wilson, W. (1975). Sentence perception as an interactive parallel process. Science, 189, 226-228.

Nevins, A. (2004) Conditions on (dis)harmony. PhD Dissertation, Department of Linguistics, MIT.

Prince, A., & Smolensky, P. (1993/2004). Optimality theory: Constraint interaction in generative grammar. Cambridge: Blackwell.

Reber, A. S. (1967). Implicit learning of artificial grammars. Journal of Verbal Learning and Verbal Behavior, 6, 855-863.

Redington, M., & Chater, N. (1996). Transfer in artificial grammar learning: A reevaluation. Journal of Experimental Psychology: General, 125(2), 123-138.

Wilson, C. (2006). Learning phonology with substantive bias: An experimental and computational study of velar palatalization. Cognitive Science, 30, 945-982.

Proceedings of the 27th West Coast Conference on Formal Linguistics

edited by Natasha Abner and Jason Bishop

Cascadilla Proceedings Project   Somerville, MA   2008

Copyright information

ISBN 978-1-57473-428-7 library binding

Ordering information

Orders for the library binding edition are handled by Cascadilla Press.
To place an order, go to www.lingref.com or contact:

Cascadilla Press, P.O. Box 440355, Somerville, MA 02144, USA
phone: 1-617-776-2370, fax: 1-617-776-2271, e-mail: sales@cascadilla.com

Web access and citation information

This entire proceedings can also be viewed on the web at www.lingref.com. Each paper has a unique document # which can be added to citations to facilitate access. The document # should not replace the full citation.

This paper can be cited as:

Finley, Sara and William Badecker. 2008. Analytic Biases for Vowel Harmony Languages. In *Proceedings of the 27th West Coast Conference on Formal Linguistics*, ed. Natasha Abner and Jason Bishop, 168-176. Somerville, MA: Cascadilla Proceedings Project.

or:

Finley, Sara and William Badecker. 2008. Analytic Biases for Vowel Harmony Languages. In *Proceedings of the 27th West Coast Conference on Formal Linguistics*, ed. Natasha Abner and Jason Bishop, 168-176. Somerville, MA: Cascadilla Proceedings Project. www.lingref.com, document #1829.