Substantive Biases for Vowel Harmony Languages

We present the results of three artificial grammar learning experiments exploring the nature of phonological representations in vowel harmony, a process whereby all vowels in a given lexical item are forced to agree in some feature value. Our results support a theory of substantively biased learning in which learners form hypotheses about novel phonological processes based on universal grammatical knowledge, either innately predetermined or a byproduct of the learning process. The substantively biased learning hypothesis predicts that learners should converge on unmarked vowel harmony patterns. Our results support this hypothesis in that learners form generalizations based on the following three properties of unmarked harmony systems: (i) learners use feature-based representations, (ii) the rules formed are general, applying to multiple affixes, and (iii) harmony patterns encode directionality in the form of stem-control and the phonetically unmarked right-to-left pattern.

While previous experiments have tested biases for naturalness in learning by comparing learning for an unnatural versus a natural rule (e.g., Peperkamp & Dupoux 2006; Pycha et al. 2003; Wilson 2003), our experiments test for learning biases using the poverty of the stimulus paradigm (Wilson, 2006), where biases are measured by performance on items that fall outside the training space. This methodology allows us to test more fine-grained naturalness patterns, and may be more immune to adult memorization strategies. In the present experiments, adult learners are exposed to stem and stem+suffix pairs where a single suffix alternates based on a vowel harmony pattern (e.g., [kinæ kinæmi] for backness harmony). Learning is then evaluated using a forced choice task, between a harmonic suffixed form and a disharmonic form (e.g., [mepemi] vs. *[mepemu]). Learning is measured in terms of the percentage of harmonic choices compared to a control condition. Critical test items contain novel segments, not heard at training. All participants were adult native speakers of English.

If languages of the world are describable in terms of abstract phonological features, learners should posit rules in terms of these features. We provide evidence for this in Experiment 1 in which participants are exposed to 4 vowels in a 6-vowel inventory, with the remaining 2 vowels held-out until test. At test, participants generalized to these novel vowels. (69% harmonic responses compared to 54% for controls for mid vowels held out of backness harmony; $p < 0.05$). Because participants were able to extend the harmony pattern outside of their training set, the results support the idea that learners make use of abstract, featural representations.

The fact that the prototypical phonological process applies irrespective of morphological context suggests that learners should form general rules that will apply in all applicable contexts, even if the context falls outside the training space (Berent et al. 2002). In Experiment 2, we exposed learners to a harmony pattern with a single suffix alternation (e.g., [-mu]-[-mi]), but participants were tested on a novel suffix alternation (e.g., [-go]-[-ge]). If learners form general rules, they should be able to accept harmonic items with a novel suffix vowel. Our results indicate that participants are able to make this generalization (76% harmonic responses compared to 47% for controls; $p < 0.01$). Together, these results support the substantively biased learning hypothesis that learning is general, involves abstract featural representations and is sensitive to universal grammatical principles.

Finally, we tested for biases on the directionality of vowel harmony. Cross-linguistically, the majority of harmony processes are either left-to-right or stem-controlled, in which the direction of harmony is determined by the stem vowel. In Experiment 3, adult speakers were trained on a color (back and round) harmony system in which front, unround vowels [i, e] triggered the prefix/suffix vowel [i] and back, round vowels [o, u] triggered the prefix/suffix vowel [u]. Participants in each hold-out condition were trained on 24 stem and stem+suffix or stem+prefix alternations (e.g., [buno, bunomu], [peki, pekimi] for suffixed training, and [buno, mubonu], [peki, mipeki] for prefixed training). Immediately following training, participants were given a forced-choice grammaticality judgment task which containing 36 items. In order to avoid generalization to novel affixes on the basis of similarity to the training affix, the novel affix contained a different consonant as the training affix (e.g., participants trained with suffix [-mi]/[-mu] were given prefix [gi-]/[gu-] and vice versa for each condition). Identity of affix consonant was counterbalanced throughout the experiment. Results indicate learning of the harmony pattern for both training conditions. Extension of harmony from stem+suffix to prefix+stem was more robust than extension of harmony from prefix+stem to stem+suffix. There were significantly fewer harmonic responses to novel suffix items compared to Old and New prefixed items in the Suffix Generalization condition, but there were equivalent responses in the Prefix Generalization condition. This suggests that learners are biased towards stem-controlled harmony and the cross-linguistically preferred right-to-left harmony, and that directionality is at least one factor at work in the learning of vowel harmony.

Our research suggests that knowledge of a given language entails knowledge of markedness that reaches beyond the processes that are found in the native language. This implies that vowel harmony makes use of general, feature-based representations that encode directionality and stem-control.