

Published in final edited form as:

Psychon Bull Rev. 2013 August ; 20(4): 780–789. doi:10.3758/s13423-013-0402-7.

Generalization to unfamiliar talkers in artificial language learning

Sara Finley

Department of Psychology, Waldorf College, 106 S 6th Street, Forest City, IA 50436

Sara Finley: sara.finley@wadorf.edu

Abstract

While there is evidence that talker-specific details are encoded in the phonetics of the lexicon (Kraljic, Samuel, & Brennan, 2008; Logan, Lively, & Pisoni, 1991) and in sentence processing (Nygaard & Pisoni, 1998), it is unclear whether categorical linguistic patterns are also represented in terms of talker-specific details. The present study provides evidence that adult learners form talker-independent representations for productive linguistic patterns. Participants were able to generalize a novel linguistic pattern to unfamiliar talkers. Learners were exposed to spoken words that conformed to a pattern in which vowels of a word agreed in place of articulation, referred to as vowel harmony. All items were presented in the voice of one single talker. Participants were tested on items that included both the familiar talker and an unfamiliar talker. Participants generalized the pattern to novel talkers when the talkers spoke with a familiar accent (Experiment 1), as well as with an unfamiliar accent (Experiment 2). Learners showed a small advantage for talker familiarity when the words were familiar, but not when the words were novel. These results are consistent with a theory of language processing in which the lexicon stores fine-grained, talker-specific phonetic details, but productive linguistic processes are subject to abstract, talker-independent representations.

Language use involves knowledge of highly specific phonetic details, as well the ability to generalize to novel situations, raising the issue of the extent to which knowledge of language relies on abstract rules versus fine-grained details. Knowledge of a specific language is general; a language user can understand almost any speaker of the language, despite the fact that every speaker has individual, idiosyncratic characteristics. Thus, the language user must be able to distinguish between speech characteristics that are idiosyncratic to the talker and speech characteristics that are shared across the language. Studying how speakers deal with unfamiliar talkers in language learning tasks can help to uncover which aspects of language processing make use of talker-specific information, and which aspects of language take place at a talker-independent level of representation.

Previous research exploring the role of talker-specific effects of language processing has focused on lexical access and phonetic patterns, with little discussion of productive, categorical linguistic patterns. This paper focuses on productive phonological patterns: systematic changes in the sounds that make up a word. For example, vowel harmony is a phonological pattern that can be found in several of the world's languages, but is not found in English. With some exceptions, Hungarian shows alternations in suffix vowels depending on the quality of the stem vowels¹. For example, the singular (dative) suffix alternates between [-nek] and [-nak], depending on the quality of the stem vowel. When the stem contains vowels pronounced in the back of the oral cavity, such as /a/ and /o/, [-nak] appears (e.g., [hajo-nak] 'ship'). When the stem contains vowels produced in the front of the oral cavity, such as /i/ and /e/, [-nek] appears (e.g., [öleles-nek] 'embrace').

¹ In the word *dogs*, *dog* is the stem and *-s* is the suffix.

In Hungarian, the formation of morphologically complex words is dependent on vowel harmony, demonstrating the interaction between phonological patterns and the lexicon. This interaction has led some researchers to propose the possibility of reducing the study of productive phonological patterns to tendencies over the lexicon (Port & Leary, 2005). Exemplar models of cognition (Goldinger, 1996, 1998; Nosofsky, 1988) serve as the basis for many of these proposals (Connine & Pinnow, 2006; Johnson, 1997; Palmeri, Goldinger, & Pisoni, 1993; Pierrehumbert, 2001; Wedel, 2006). Exemplar models of language are supported by the robust finding that talker familiarity serves as an aid to lexical access (Connine & Pinnow, 2006; Creel, Aslin, & Tanenhaus, 2008; Creel & Tumlin, 2009; Goldinger, 1998; Palmeri, et al., 1993; Pisoni, 1997). In these studies, listeners were faster and more accurate when recalling words spoken by a familiar talker than an unfamiliar talker. These results suggest that talker-specific information is encoded in the lexicon. In addition, each token of a word is stored in memory along with the fine-grained phonetic characteristics of that token.

Because lexical entries encode talker specific information, there is reason to believe that phonological patterns may also make use of talker-specific information (Pierrehumbert, 2001). While no studies have focused on categorical phonological patterns like vowel harmony, research has tested learners' ability to generalize to novel talkers when learning a novel phonetic contrast, such as the difference between /l/ and /r/ that is present in English but not Japanese. Japanese learners of English were able to extend the novel phonetic category to novel talkers only when participants were trained on stimuli that included multiple talkers (Lively, Pisoni, & Logan, 1992; Logan, et al., 1991; Magnuson et al., 1995). This suggests that phonetic contrasts may be formed via talker-specific representations. However, categorical phonological patterns like vowel harmony may be represented differently than phonetic contrasts that tend to make greater use of fine-grained phonetic details. It is therefore unclear whether categorical phonological patterns like vowel harmony will show the same talker-specific effects.

There is evidence that talker-specific knowledge is used in abstract, sentence-level processing (Nygaard & Pisoni, 1998). Participants were trained on talker identities by listening to several talkers produce full sentences. Participants showed better recognition for individual words in test sentences that were spoken by familiar talkers, demonstrating that talker-specific information is used in high-level language processing. These results support the view that linguistic processes, including phonological patterns, may be subject to talker-specific effects. Because Nygaard and Pisoni (1998) only tested word recognition, it is unclear whether talker-specific processing extends to processing novel stimuli from categorical phonological patterns like vowel harmony, when participants are trained on a single phonological pattern spoken by one individual talker.

In addition, phonetic differences between categories of sounds may make some sounds more applicable to generalization to novel talkers. In a perceptual learning task, Kraljic and Samuel (2006) showed that participants can extend a novel category boundary to unfamiliar talkers, but that this generalization is constrained by the specific sounds involved in the contrast (Kraljic & Samuel, 2005, 2007), as well as the particular behavior of the talker (Kraljic, et al., 2008). Speakers were less likely to extend the novel accent to an unfamiliar talker if the familiar talker spoke with a pen in the mouth (Kraljic, et al., 2008).

This paper uses an artificial grammar learning paradigm to study the effects of talker familiarity on processing novel phonological patterns. Previous artificial language learning experiments have shown that adults can learn a novel vowel harmony pattern after brief exposure (Finley & Badecker, 2008, 2009a, 2009b, 2010; Koo & Cole, 2006; Moreton, 2008; Pycha, Nowak, Shin, & Shosted, 2003; Skoruppa & Peperkamp, 2011). In addition,

learners of novel phonological patterns show robust generalization to novel items. For example, participants in a study by Finley and Badecker (2009a) heard a novel back/round vowel harmony pattern that contained four vowels from a six-vowel inventory. Following exposure, participants were given a two-alternative forced-choice task in which participants chose between a harmonic form and a disharmonic form. For example, participants chose between harmonic [bodumu] and disharmonic *[bodumi]². Test items were divided into three categories: Old Items, which were words that appeared in training, New Items, which were words that did not appear in training, but contained the same vowels and consonants as the training set, and New Vowel Items, which were words that contained vowels not present in the training set. Participants extended the vowel harmony pattern to novel vowels, suggesting that novel phonological processes are learned in terms of abstract representations.

Previous language learning experiments, specifically those exploring phonological patterns, made use of a single talker during exposure. These studies did not test for generalization to novel talkers. This leaves open the possibility that learners infer talker-specific representations rather than language-specific representations. Some insight to how learners respond to novel talkers in an artificial language learning task may be gained through examination of artificial grammar learning experiments that explored the role of transfer across modalities, such as from a spoken pattern to an analogous written pattern (Dienes & Altmann, 1997). These studies found robust transfer across modalities, but also found a transfer deficit; correct responses decreased in the novel modality. While transfer across modalities is different from transfer across talkers, understanding the level of representation for which items are stored tests the limits of human generalization. These findings are useful in creating theories of the levels of representation for both linguistic and nonlinguistic pattern learning. In addition, the existence of a 'transfer deficit' in generalization across modalities suggests that learners will show deficits across talkers as well (Dienes & Altmann, 1997).

The present study tests for the existence of talker-independent representations in novel phonological patterns. Evidence for talker-independent representations will be found if participants are able to extend the novel harmony pattern from a familiar talker to an unfamiliar talker. This extension will be measured in two ways. First, if correct responses to the unfamiliar talker exceed correct responses in the Control condition, it suggests that participants have a representation of the vowel harmony pattern that goes beyond familiar talker. Second, transfer deficits will be assessed by comparing responses to the unfamiliar talker to corresponding responses to the familiar talker. Statistically significant differences between familiar and unfamiliar talkers provide evidence for talker-specific details within the representation of the harmony pattern. A division between Old Items (words heard in training) and New Items (novel words not heard in training) allows for comparisons to be made with respect to lexical familiarity. These different comparisons yield several possible outcomes and interpretations. The four most probable outcomes are listed in Table 1.

If participants in the experimental condition fail to extend the pattern to unfamiliar talkers, as compared to the Control condition, it suggests that the harmony pattern was learned using talker-specific representations. If there is a failure to extend the harmony pattern to unfamiliar talkers, one expects that there will also be significant transfer deficits for both Old and New items. Because failure to extend to unfamiliar talkers has been shown in previous studies, (Kraljic & Samuel, 2007), it is possible that learners in the present study will also fail to extend the harmony pattern to unfamiliar talkers.

² The '*' indicates an ungrammatical, disharmonic item.

If, in addition to storing information about the familiar talker, speakers have access to a general, talker-independent representation of the pattern, learners may show a significant generalization to unfamiliar talkers compared to the Control condition, but with transfer deficits. Under this pattern of results, there are two possible outcomes: transfer deficits for both Old and New Items or transfer deficits for Old Items only. If transfer deficits occur for all items, it suggests that learners make use of talker-specific details to learn the harmony pattern, but when exposed to the same pattern spoken in an unfamiliar voice, the learner must extend the stored representations that contain the familiar talker to the unfamiliar talker, resulting in a transfer deficit for all items. However, it is also possible that transfer deficits will only occur for Old Items. Much of the research supporting talker-specific representations has focused on lexical access. There is evidence that words are easier to access if they are spoken by a familiar talker (Nygaard, Sommers, & Pisoni, 1994; Palmeri, et al., 1993). If the general phonological pattern is learned under talker-independent mechanisms, but the lexicon incorporates talker-specific information, one should expect that participants will correctly respond to items spoken with an unfamiliar talker (compared to a Control condition), but will show transfer deficits for Old items only. New Items will not show transfer deficits because the learner has no prior representation of these words in the lexicon.

The final possible outcome assumes that phonological patterns are stored without any talker-specific information. In this case, participants will extend the harmony pattern to unfamiliar talkers without any transfer deficit. This pattern of results would support the strongest version of talker-independence. In order for this outcome to occur, learners must show the same pattern of results for the familiar talker and the unfamiliar talker, even for familiar (Old) test items. This possible pattern of results supports a view in which both the lexicon and the phonological pattern make use of talker-independent representations.

The experiments in the present paper explore the role of talker independence in learning a novel phonological pattern. The ability to extend a novel vowel harmony pattern to an unfamiliar talker will help to shed light on the nature of learning and representations of phonological knowledge.

EXPERIMENT 1

Method

Participants—All participants were adult native speakers of English with no knowledge of a vowel harmony language. Seventy-two University of Rochester undergraduate students and affiliates were paid \$10 for their participation. There were 40 participants in the critical conditions and 32 participants in the No-Training Control condition.

Design—An artificial grammar learning task was used to assess the ability to generalize novel phonological patterns to novel talkers. In an artificial grammar task, a novel language is created that has a specific pattern or characteristic, such as vowel harmony. Participants in the experimental conditions were exposed to a vowel harmony pattern spoken by a single talker, either male or female. The training phase was followed immediately by a two-alternative forced-choice test that contained items spoken by the familiar talker as well as a novel talker. The general design of the experiment was based on Finley and Badecker (2009a), described above. All phases of the experiment were presented using PsyScopeX (Cohen, MacWhinney, Flatt, & Provost, 1993).

Participants in the critical conditions were assigned to either the Male Talker Training condition or the Female Talker Training condition. Tokens in both conditions were identical,

except that tokens in the Male Talker Training condition were spoken by a male talker, while tokens in the Female Talker Training condition were spoken by a female talker.

Training items in the critical conditions were of the form of a bisyllabic ‘stem’ word (of the form CVCVC, where C is a consonant and V is a vowel) immediately followed by a back/round harmonic ‘suffixed’ word (of the form CVCVC-V), but without any meanings associated with the items. For example, the stem [gemit] was followed by the harmonic [gemit-e]. Stems contained front vowels [i, e] or back vowels [o, u]. The suffix was a vowel that alternated between [-e] and [-o]. The suffix [-e] appeared when the stem vowels were front (e.g., [mekin, mekin-e]); the suffix [-o] appeared when the stem vowel was back (e.g., [poduk, poduk-o]). Examples of training stimuli can be found in Table 2; full stimuli lists are in the appendix.

All stimuli contained the same consonant inventory, [p, b, t, d, k, g, m, n] and the same vowel inventory, [i, u, e, o]. Sixteen training items were created for each critical condition. The training stimuli were counterbalanced to contain all possible combinations of vowel sounds. Suffixed words were produced semi-randomly with the condition that all stimuli were not homophonous with an English word. The final profile of the stimuli was counterbalanced to appropriately contain equivalent use of the different vowel patterns in the stems.

Training was followed by a two-alternative forced-choice test with 40 test items. Each test item contained two suffixed forms, one harmonic and one disharmonic. For example, participants chose between harmonic [gemit-e] and disharmonic *[gemit-o]. Half of the test items were presented in the female voice, while the other half of test items contained the male voice. One group of participants heard all 40 items in an unblocked, random order, with familiar and unfamiliar talkers mixed in a random order (n = 16). All other participants (n = 24) heard test items presented in two blocks of 20 items, an all-male and an all-female block, with the order of presentation counterbalanced such that half of participants heard the male test items first³. Each block contained 10 words that had appeared in the exposure phase and 10 words that did not appear in exposure phase; these items were presented in a random order. This amounted to four total test conditions: Old Items (Familiar Talker), New Items (Familiar Talker), Old Items (Unfamiliar Talker) and New Items (Unfamiliar Talker).

Thirty-two participants were assigned to a No-Training Control condition. The control condition was designed to ensure that all effects were due to learning, as opposed to an inherent response bias. In this No-Training Control condition, participants received the same test items as participants in the critical conditions. Half of the Control participants responded to items in an unblocked, random order, while the other half responded to items in blocks. Of these participants, eight responded to female items first, eight participants responded to male items first, and 16 participants responded to male and female items in a random order. While the Control participants had not heard any of the test items during the exposure phase and were thus all ‘new’, the test items were matched to the appropriate test condition based on the critical condition⁴.

As noted in Table 1 and the description of the table, the extent to which learning novel, categorical phonological patterns is based in talker-independent representations should be

³ Results did not indicate any difference in responses depending on blocked or random order of presentation. The overall average response rate was 0.78 for random and 0.76 for blocked.

⁴ Finley and Badecker (2009a) used a ‘stem’ only as well as a mono-syllabic control condition. In these conditions, participants are given non-harmonic (neither harmonic or disharmonic) pattern to learn. Responses to these control conditions were also close to chance, making it unlikely that the Control condition in the present experiment was any better or worse than previous control conditions.

reflected in the extent to which learners are able to extend the harmony pattern to an unfamiliar talker, both within comparison to the Control condition, as well as compared to responses to a familiar talker, as a transfer deficit.

Materials

The naturally produced stimuli were recorded in a sound attenuated booth with a 22kHz sampling rate from two native speakers of American English, one male and one female. Both speakers spent the majority of their childhood in the same region of the United States, Upstate New York. While the speakers had no knowledge of the specifics of the experimental design, they were aware that the items would be used in an artificial language learning task. All stimuli were phonetically transcribed, and presented to the speakers in written format. The speakers were instructed to produce all vowels as clearly and accurately as possible, even in unstressed positions. Stress was produced on the first syllable in all forms. Because talkers were told to speak naturally, length of utterances was not controlled for. Thus, there were differences in durations between the male and the female talkers, with the female talker being slightly longer. Male suffixed items averaged 835ms, with a range of 760–1000ms. Female suffixed items averaged 1303ms, with a range of 881–1401ms. Such differences in length make the talkers even more distinguished, which, if anything, should prevent generalization to the novel talker. All items were scaled to the same intensity level. All sound files were edited in Praat (Boersma & Weenink, 2005).

Procedure

All phases of the experiments took place at Macintosh computer with stimuli presented using over-ear headphones. Participants were allowed to adjust the volume of the headphones to a comfortable level. Participants received both written and verbal instructions. Participants in the critical conditions were told that they would be listening to words from a language they had never heard before, but that they need not memorize the forms. Following the exposure phase, participants in the critical conditions were told that they would hear two words, one from the language they just heard, the other not from the language; their job was to select the word from the language. If they believed that the first word was from the language, they were to press the 'a' key on the keyboard, and if they believed that the second word was from the language, they were to press the 'l' key on the keyboard. They were told to respond as quickly and accurately as possible, but to wait until they heard both possibilities before responding. The experiment took approximately 15 minutes.

Participants in the Control condition were only given the test items, and were therefore not given any exposure items to compare to during testing. This means that the directions given to participants in the critical conditions were not appropriate for the Control condition. Instead, participants were told that they would be making judgments about pairs of words. Their task was to decide which of two words they preferred, based on any criteria they chose, and that there was no 'correct' or 'incorrect' response. Participants in the Control condition were given the same set of test items, and responded using the computer keyboard in the same manner as participants in the Critical conditions. The control experiment took approximately 5 minutes to complete.

Results

Proportions of same language, harmonic responses are provided in Table 3. Results are categorized in terms of the Control condition and the two Critical Conditions: Male Talker Training and Female Talker Training, and divided by Old Items and New Items for Familiar and Unfamiliar Talker Test Items. Control Items have male as the default Familiar Talker, but statistical comparisons were made according to the appropriate gender.

A 2 (Training Gender) x 2 (Talker Familiarity) x 2 (Lexical Familiarity) mixed-design ANOVA was used to compare responses in the Female Talker Training Condition to the Male Talker Training Condition. This comparison ensured that responses did not differ based on the gender of the talker heard during the exposure phase. There was no effect of gender, $F(1, 38) = 1.75, p = 0.19, \eta^2 = 0.044$, suggesting that both training conditions were equally able to learn the harmony pattern. There was a significant effect of Lexical Familiarity, $F(1, 38) = 29.83, p < 0.0001, \eta^2 = 0.44$, and no interaction by Gender, $F < 1$, suggesting that participants were more accurate with Old Items than New Items. There was a marginally significant effect of Talker Familiarity, $F(1, 38) = 3.10, p = 0.084, \eta^2 = 0.075$, with no interaction by gender, $F < 1$, suggesting that participants were more accurate on items heard in the familiar voice. There was no interaction between Lexical Familiarity and Talker Familiarity, $F < 1$, and no three-way interaction, $F < 1$.

Responses to the unfamiliar talker in the Critical conditions were compared to the corresponding items in the Control condition. In order to match items in the Critical and the Control conditions, two separate comparisons were performed: Female Talker Training vs. Control, and Male Talker Training vs. Control, via Bonferroni corrected independent sample t-tests. There were significantly more harmonic responses for all test conditions. There were significantly more harmonic responses for the Male Talker Training condition compared to the Control condition for New-Unfamiliar Talker Items, $t(50) = 7.94, p < 0.0001$, and Old-Unfamiliar Talker Items, $t(50) = 5.44, p < 0.0001$, as well as for the comparisons between the Female Talker condition and the Control condition: Old-Unfamiliar Talker Items, $t(50) = 6.56, p < 0.0001$ and New-Unfamiliar Talker Items, $t(50) = 3.72, p < 0.001$.

Transfer deficits were detected using planned comparisons between Familiar and Unfamiliar talkers for Old and New Items. Male and Female Talker Training conditions were combined to increase power, as there was no difference or interaction by gender in the ANOVA. Means and standard errors are presented in Figure 1. There was a marginally significant difference between Familiar and Unfamiliar Talkers for Old Items, $t(39) = 1.96, p = 0.057$ (0.85 vs. 0.80), but there was no significant difference for New Items, $t(39) = 0.82, p = 0.42$ (0.73 vs. 0.71). These results are consistent with the hypothesis that talker-specific effects are strongest for known words.

These results suggest that learning a productive phonological pattern takes place at a level at which talker identity does not impede judgments made outside of the lexicon. There was a marginal transfer deficit for Old Items, but no effect for New Items⁵, suggesting that any talker-specific representations are delegated to the lexicon.

Discussion

Learners in the present experiment generalized the harmony pattern to novel words spoken by novel talkers. Previous research (Houston & Jusczyk, 2000) demonstrated that infants are more likely to generalize to a novel talker if they first hear either a familiar word or a familiar talker. It is possible that learners generalized to the novel talker for novel items only because they heard familiar items first. To verify this, we partitioned all New/Unfamiliar Talker test items that occurred before any other test item. While many participants did not hear this type of test item first (i.e., those that heard the familiar talker in the first block), all responses to items in this category were correct (i.e., listeners chose the harmonic response on all trials).⁶

⁵ Analyses were also conducted based on reaction time, with similar results. Because the task was not a speeded judgment task, and involved a binary response, any analysis on reaction time must be taken with extreme caution, and are thus not included in the main text.

Another factor that may have led to increased generalization to unfamiliar talkers was the instructions given to participants at the time of training. The instructions in Experiment 1 noted that participants ‘need not memorize’ the words that they heard. This instruction may have primed participants to form an abstract phonological rule. To test this possibility, we ran a small set of participants ($n=10$) with the instructions ‘please memorize the words you hear’. The overall pattern of results was consistent with the results of Experiment 1.

Experiment 1 demonstrates that adult learners are able to extend a novel phonological pattern to unfamiliar talkers. However, all the talkers in Experiment 1 spoke with an American English accent. It is possible that formation of talker-independent representations is contingent on familiarity with the accent in question. For example, speakers may know where to look for talker-specific effects versus talker-independent effects in English speech, and were able to extend the pattern to unfamiliar talkers based on this prior knowledge. Experiment 2 explores this possibility through a replication of Experiment 1, using Turkish talkers and English-speaking participants.

EXPERIMENT 2

Experiment 2 was identical to Experiment 1 except that the stimuli made use of Turkish-speaking talkers, rather than English-speaking talkers. Replacing English-speaking talkers with Turkish-speaking talkers increases the ecological validity of the present experiment because learning a novel language requires learning a novel accent. Because Turkish naturally has vowel harmony, speakers hearing harmony in a Turkish dialect are more closely simulating the experience of a second language learner learning a novel vowel harmony pattern.

Method

Participants—Participants were adult native speakers of English with no knowledge of a vowel harmony language, and had not participated in a previous vowel harmony learning experiment. Forty University of Rochester undergraduate students and affiliates were paid \$5–10 for their participation. There were 28 participants in the critical condition and 12 participants in the No-Training Control condition.

Design—The design of Experiment 2 was identical to Experiment 1 except that the talkers used in the experiment were Turkish speakers. Half of the participants were exposed to the female Turkish talker, and the other half were exposed to the male Turkish talker. All participants heard all test items in a random, mixed order. Because the blocked design did not vary compared to the random, for simplicity, all participants received the same design.

Materials

Materials were collected in the same manner as Experiment 1 with a few minor differences. First, two native Turkish speakers from Istanbul recorded the experimental stimuli. All of the vowels and consonants from Experiment 1 are found in Turkish, making it possible to use identical stimuli. However, all languages differ in their pronunciation of vowels and consonants, meaning that the Turkish stimuli were qualitatively different from the English stimuli. For example, English vowels are often produced as diphthongs, while Turkish vowels are not. While neither of the talkers were native English speakers, the talkers were fluent in English. Second, the talkers were told to produce the words to represent words

⁶ Because there were a relatively small number of test items (10), we ran a small set of participants ($n = 4$) who heard all 10 New/Unfamiliar Talker items before any other item. Each of the four participants showed results that were consistent with the results of Experiment 1.

spoken in Turkish as closely as possible. In order to create a natural environment for the Turkish speakers, the talkers were asked to ‘speak as clearly and accurately as possible’, making it difficult to control for length of utterances. Thus, there were differences in durations between the male and the female talkers (with the female talker being slightly longer). Male suffixed items averaged 590ms, with a range of 463–810ms⁷. Female suffixed items averaged 776ms, with a range of 628–1033ms. These differences in length make the talkers even more distinguished. If anything, this difference should prevent generalization to the novel talker.

Procedure

The procedure was identical to that of Experiment 1.

Results

Proportions of same language, harmonic responses are provided in Table 3. Results are categorized in terms of the Control condition and the two Critical Conditions: Male Talker Training and Female Talker Training, and divided by Old Items and New Items for Familiar and Unfamiliar Talker Test Items. Control Items have male as the default Familiar Talker, but statistical comparisons were made according to the appropriate gender.

A mixed-design ANOVA was used to compare responses in the Female Talker Training Condition to the Male Talker Training Condition using a 2 (Gender) x 2 (Talker Familiarity) x 2 (Lexical Familiarity) design. There was no effect of gender, $F(1, 26) = 1.26, p = 0.27, \eta^2 = 0.046$, suggesting that participants learned the harmony pattern in both experimental conditions. As in Experiment 1, there was a significant effect of Lexical Familiarity, $F(1, 26) = 6.87, p < 0.05, \eta^2 = 0.21$, and no interaction by Gender, $F < 1$, suggesting that participants were more accurate with Old Items compared to New Items. There was no effect of Talker Familiarity, $F < 1$, with no interaction by gender, $F < 1$, suggesting that participants were not significantly more accurate on items heard in the familiar voice. This may be due to the fact that participants in both the Male Talker Training condition and Female Talker Training condition were more accurate on Old Items when they were heard in a familiar voice, but were numerically more accurate on New Items when they appeared in an unfamiliar voice. This is reflected in the marginal interaction between Lexical Familiarity and Talker Familiarity, $F(1, 26) = 3.05, p = 0.092, \eta^2 = 0.10$.

Responses to the unfamiliar talker in the Critical conditions were compared to the corresponding items in the Control condition. In order to match items in the Critical and the Control conditions, two separate comparisons were performed: Female Talker Training vs. Control and Male Talker Training vs. Control via Bonferroni corrected independent sample t-tests. There were significantly more harmonic responses for the Male Talker Training condition compared to the Control condition for New-Unfamiliar Items, $t(26) = 4.47, p < 0.05$, and Old-Unfamiliar Items $t(26) = 2.47, p < 0.05$, as well as for the comparisons between the Female Talker condition and the Control condition: Old-Unfamiliar, $t(22) = 5.64, p < 0.001$ and New-Unfamiliar, $t(22) = 2.61, p < 0.05$. These results suggest that participants in the Critical conditions learned the harmony pattern and extended that pattern to unfamiliar talkers.

Transfer deficits were detected using planned comparisons between Familiar and Unfamiliar talkers for Old and New Items. Male and Female Talker Training conditions were combined to increase power, as there was no difference or interaction by gender in the ANOVA.

⁷ It is unclear whether the differences in speech rate are due to differences between English and Turkish or due to the fact that the talkers in Experiment 1 were more experienced in producing experimental stimuli, and thus spoke more carefully.

Means and standard errors are presented in Figure 1. There was a marginally significant difference between Familiar and Unfamiliar Talkers for Old Items, $t(27) = 1.71, p = 0.099$ (0.79 vs. 0.71), but there was no difference between Familiar and Unfamiliar Talker items for New Items, $t(27) = 1.56, p = 0.13$ (0.64 vs. 0.68), and this trend was in the opposite direction of a talker-specific interpretation. These results are consistent in part with the hypothesis that productive linguistic patterns make use of talker-independent representations for novel items, but lexical representations are more likely to make use of fine-grained talker-specific representations.

The results of Experiment 2 parallel those of Experiment 1. This suggests that the ability to form a general phonological pattern is not contingent on familiarity with the accent of the talker.

Discussion

Overall, the responses in Experiment 2 were slightly less accurate than responses in Experiment 1 (e.g., the overall average accuracy in Experiment 1 was 77%, compared with 70.5% in Experiment 2, which was marginally significant $F(1, 66) = 3.08, p = 0.0084$). It is likely that this decline in accuracy is due to the fact that learners in Experiment 2 had to cope with a novel accent in the training phase.

GENERAL DISCUSSION

Participants in the present study learned a novel back/round vowel harmony pattern, and extended that harmony pattern to an unfamiliar talker. This occurred for both Turkish and English speaking talkers. Marginal transfer deficits occurred for Old Items, but no deficits appeared for New items, suggesting that talker-independent mechanisms are at work when listeners make novel judgments regarding categorical phonological patterns.

In the present experiments, the talkers shared the same dialect. It is possible that generalization of a novel pattern across speakers of different dialects may be less robust than generalization across the same dialect. This is consistent with the phonetic relevance hypothesis (Sommers & Barcroft, 2006) that states that the elements relevant to the phonetic code are the ones that listeners will apply to multiple talkers. In natural learning situations, children are exposed to multiple talkers, but will typically be exposed to a single primary talker, such as the primary caregiver. Young infants are able to generalize properties of speech from one speaker to another, so long as the talkers' voices are relatively similar (Houston & Jusczyk, 2000; Schmale & Seidl, 2009). Schmale and Seidl (2009) showed that 9-months old infants can accommodate different speakers from their native dialect, but could no longer do so when the novel talker spoke in an unfamiliar dialect. One goal for future research is to explore how the features of multiple talkers during exposure affect extension to unfamiliar talkers at test.

The vowel harmony pattern in the present study was relatively simple, including only four vowels. This raises the possibility that learners did not form a productive harmony pattern, but a simple association between stem and suffix vowels. This interpretation is unlikely because prior research has demonstrated that learners of a novel vowel harmony pattern are able to generalize to novel vowels outside a four-vowel inventory (Finley & Badecker, 2009a). Further, learning a novel vowel harmony pattern decreases when the associations between stem and suffix vowels are arbitrary (Pycha et al., 2003). Finally, the categorical nature of the pattern used in the present study holds despite its simplicity.

The present study extends prior research demonstrating that novel phonological patterns can be learned very rapidly, as participants were only given about 10 minutes of exposure. This

raises the possibility that learners did not have enough time to learn the talker's idiosyncrasies, resulting in minimal transfer deficits. While this is an important possibility, the goal of the present study was to assess whether learning categorical phonological patterns requires talker-specific representations. Because the harmony pattern was successfully learned in a matter of minutes, either learning does not require a strong sense of familiarity with the talker, or the idiosyncrasies of the talker were learned rapidly. Learners may have become familiar with the talker very quickly because only a single talker was heard during training. It is unlikely, however, that the use of a single talker lead the learner to assume that all aspects of the talker's speech were general to the language at hand. Previous studies have shown that generalization to novel talkers increases with the number of talkers heard during training (Magnuson, et al., 1995).

Categorical, exceptionless patterns may be more susceptible to talker-independent representations than fine-grained phonetic patterns (Lively, et al., 1992; Logan, et al., 1991; Magnuson, et al., 1995). Phonetic patterns tend to be more continuous and variable in terms of rate and consistency of application. Phonological patterns tend to be described in terms of categorical features and segments, and exhibit lower levels of variability. While categorical linguistic patterns are subject to exceptions, these exceptions tend to be principled (Zondefeld, 1978). An important question for future research would be to understand the cases, if any, where talker-specific details play a role in judgments for novel instances of a categorical phonological pattern. One possibility is that learners rely more on talker-specific details when the phonological pattern shows high levels of exceptions or lexical constraints. Another possibility is that using a task that orients the learner towards the phonetic details of the pattern, such as the tasks used in previous studies (Magnuson & Yamada, 1995; Nygaard & Pisoni, 1998) may yield more talker-specific responses. One goal of future research is to determine both the properties of the pattern, as well as the properties of the tasks that create an environment where learners are more prone to respond with respect to talker-specific details.

There is evidence that talker-specific details found in phonetic patterns are stored in the lexicon (Goldinger, 1998) and are available during lexical access (Salverda et al., 2007). Lexical items may be subject to talker-specific details even for words that are learned in a short amount of time and have no semantic content. This predicts that if the vowel harmony pattern in the present study were subject to greater phonetic or lexical variability, generalization to unfamiliar talkers may decline. Learners in the present experiments showed a marginally significant transfer deficit for Old Items, but not New items. This supports the hypothesis that learning novel categorical patterns involves multiple levels of representations (Luce, Goldinger, Auer, & Vitevitch, 2000; Luce & Large, 2001). Understanding how talker-specific details are used at various levels of representation may shed light on the mechanisms required to integrate productive patterns into an exemplar model of cognition (Goldinger, 1996, 1998; Nosofsky, 1988). Creating a theory that allows for various talker-specific effects at the phonetic level, the lexical level, as well as the categorical level, may lead the way to a better understanding of the interaction between low level speech processes, categorical phonological patterns, and the lexicon (Pierrehumbert, 2001, 2003).

While talker-specific effects are clearly language specific, there are important parallels to other areas of cognitive science. The study of learning and generalization across novel items has important consequences for understanding the mechanisms that underlie learning and generalization (Dienes & Altmann, 1997). These consequences apply to both linguistic and non-linguistic cognitive processes. For example, object recognition and category discrimination requires the ability to distinguish between individual exemplars, but also the ability to generalize to novel instances (Nosofsky, 1988). Understanding how different

patterns in cognition are subject to different levels of representation and generalization may help to create a unified theory of cognition.

The present study tested the role of talker-specific representations in learning novel phonological patterns. Participants were exposed to a novel vowel harmony pattern in the voice of a single talker, and then tested on both a familiar and an unfamiliar talker. Participants generalized to novel talkers, supporting the hypothesis that learners make use of abstract representations in making judgments regarding novel categorical phonological patterns. Learners showed a marginal transfer deficit for Old Items, but no deficits for New items, supporting a theory that learning mechanisms for phonological patterns make use of representations that allow for generalization beyond the familiar.

References

- Boersma P, Weenink. Praat: Doing phonetics by computer. 2005
- 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.
- Connine C, Pinnow E. Phonological variation in spoken word recognition: Episodes and abstractions. *The Linguistic Review*. 2006; 23:235–245.
- Creel SC, Aslin RN, Tanenhaus MK. Heeding the voice of experience: The role of talker variation in lexical access. *Cognition*. 2008; 106:633–644. [PubMed: 17507006]
- Creel, SC.; Tumlin, MA. Talker variability is intrinsic to word representations: Evidence from on-line processing of spoken words. In: Taatgen, NA.; van Rijn, H., editors. *Proceedings of the 31st annual Cognitive Science Conference*. Austin, TX: Cognitive Science Society; 2009. p. 845-850.
- Dienes, Z.; Altmann, G. Transfer of implicit knowledge across domains? How implicit and how abstract?. In: Berry, D., editor. *How implicit is implicit learning?*. Oxford: Oxford University Press; 1997. p. 107-123.
- 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. Linguistic and non-linguistic influences on learning biases for vowel harmony. In: Ohlsson, S.; Catrambone, R., editors. *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society; 2010. p. 706-711.
- Goldinger SD. Words and voices: Episodic traces in spoken word identification and recognition memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*. 1996; 22:1166–1183.
- Goldinger SD. Echoes of echoes? An episodic theory of lexical access. *Psychological Review*. 1998; 105(2):251–279.
- Houston DM, Jusczyk PW. The role of talker-specific information in word segmentation by infants. *Journal of Experimental Psychology: Human Perception and Performance*. 2000; 26(5):1570–1582. [PubMed: 11039485]
- Johnson, K. Speech perception without speaker normalization. In: Johnson, K.; Mullenix, JW., editors. *Talker variability in speech processing*. San Diego: Academic Press; 1997. p. 145-165.
- Koo, H.; Cole, J. In: Botinis, A., editor. *On learnability and naturalness as constraints on phonological grammar*; *Proceedings of ISCA Tutorial and Research Workshop on Experimental Linguistics*; Athens. 2006. p. 174-177.
- Kraljic T, Samuel A. Perceptual learning in speech: Is there a return to normal? *Cognitive Psychology*. 2005; 51:141–178. [PubMed: 16095588]
- Kraljic T, Samuel A. Generalization in perceptual learning of speech. *Psychonomic Bulletin & Review*. 2006; 13:262–268. [PubMed: 16892992]

- Kraljic T, Samuel A. Perceptual adjustments to multiple speakers. *Journal of Memory and Language*. 2007; 56:1–15.
- Kraljic T, Samuel A, Brennan S. First impressions and last resorts: How listeners adjust to speaker variability. *Psychological Science*. 2008; 19(4):332–228. [PubMed: 18399885]
- Lively, SE.; Pisoni, DB.; Logan, JS. Some effects of training Japanese listeners to identify English /r/ and /l/. In: Tohkura, Yi; Vatikiotis-Bateson, E.; Sagisaka, Y., editors. *Speech perception, production and linguistic structure*. Burke, VA: IOS Press; 1992. p. 175-196.
- Logan JS, Lively SE, Pisoni DB. Training Japanese listeners to identify English /r/ and /l/: A first report. *Journal of the Acoustical Society of America*. 1991; 89(2):874–886. [PubMed: 2016438]
- Luce PA, Goldinger SD, Auer ET, Vitevitch MS. Phonetic priming, neighborhood activation, and PARSYN. *Perception and Psychophysics*. 2000; 62:615–625. [PubMed: 10909252]
- Luce PA, Large NR. Phonotactis, density, and entropy in spoken word recognition. *Language and Cognitive Processes*. 2001; 16(5/6):565–581.
- Magnuson, JS.; Yamada, RA. The effects of talker variability on the acquisition of non-native speech contrasts. *Proceedings of the 1995 International Congress of Phonetic Sciences*; 1995. p. 306-309.
- Magnuson, JS.; Yamada, RA.; Tohkura, Yi; Bradlow, AR.; Lively, SE.; Pisoni, DB. The role of talker variability in non-native phoneme training. *Proceedings of the 1995 Spring Meeting of the Acoustical Society of Japan*; 1995. p. 393-394.
- Moreton E. Analytic bias and phonological typology. *Phonology*. 2008; 25:83–127.
- Nosofsky R. Exemplar-based accounts of relations between classification, recognition and typicality. *Journal of Experimental Psychology: Learning, Memory and Cognition*. 1988; 14(4):700–708.
- Nygaard LC, Pisoni DB. Talker-specific learning in speech perception. *Perception & Psychophysics*. 1998; 60(3):355–376. [PubMed: 9599989]
- Nygaard LC, Sommers MS, Pisoni DB. Speech perception as a talker-contingent process. *Psychological Science*. 1994; 5(1):42–46. [PubMed: 21526138]
- Palmeri TJ, Goldinger SD, Pisoni DB. Episodic encoding of voice attributes and recognition memory for spoken words. *Journal of Experimental Psychology: Learning, Memory and Cognition*. 1993; 19(2):309–328.
- Pierrehumbert, J. Exemplar dynamics: Word frequency, lenition and contrast. In: Bybee, JL.; Hopper, P., editors. *Frequency effects and emergent grammar*. Amsterdam: John Benjamins; 2001. p. 137-157.
- Pierrehumbert, J. Probabilistic phonology: Discrimination and robustness. In: Bod, R.; Hay, J.; Jannedy, S., editors. *Probabilistic linguistics*. Cambridge, MA: The MIT Press; 2003. p. 177-228.
- Pisoni, DB. Some thoughts on ‘normalization’ in speech perception. In: Johnson, K.; Mullenix, J., editors. *Talker variability in speech perception*. San Diego: Academic Press; 1997. p. 9-32.
- Port RF, Leary AP. Against formal phonology. *Language*. 2005; 81:927–964.
- 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.
- Salverda AP, Dahan D, Tanenhaus MK, Crosswhite K, Masharov M, McDonough J. Effects of prosodically modulated sub-phonetic variation on lexical competition. *Cognition*. 2007; 105:466–476. [PubMed: 17141751]
- Schmale R, Seidl A. Accommodating variability in voice and foreign accent: Flexibility of early word representations. *Developmental Science*. 2009; 12(4):583–601. [PubMed: 19635085]
- Skoruppa K, Peperkamp S. Adaptation to novel accents: Feature-based learning in context-sensitive phonological regularities. *Cognitive Science*. 2011; 35(2):348–366. [PubMed: 21429003]
- Sommers M, Barcroft J. Stimulus variability and the phonetic relevance hypothesis: Effects of variability in speaking style, fundamental frequency, and speaking rate on spoken word identification. *Journal of the Acoustical Society of America*. 2006; 119(4):2406–2416. [PubMed: 16642853]
- Wedel A. Exemplar models and language change. *The Linguistic Review*. 2006; 24:147–185.
- Zonnefeld, W. *A Formal Theory of Exceptions in Generative Phonology*: Walter de Gruyter. 1978.

Appendix: Full Stimuli Lists

Training Items	Old Test Items	New Word Test Items
budok, budoko	degibe, *degibo	*bipido, bipide
digib, digibe	*giteko, giteke	*mukobe, mukobo
dupob, dupobo	*budoke, budoko	tidipe, *tidipo
gemit, gemite	*dupobe, dupobo	*toguke, toguko
gitek, giteke	gemite, *gemito	*bidito, bidite
kimet, kimete	kukopo, *kukope	butoko, *butoke
kukop, kukopo	monuko, *monuke	godomo, *godome
mekin, mekine	netepe, *netepo	kukogo, *kukoge
midik, midike	*nopube, nopubo	nopuko, *nopuke
monuk, monuko	*tikepo, tikepe	*pedebo, pedebe
netep, netepe		
nopub, nopubo		
puduk, puduko		
tikep, tikepe		
todup, todupo		
tokot, tokoto		




Figure 1. Transfer Deficits for Experiments 1 and 2 (Means and Standard Errors)

Table 1

Possible Results and Interpretations

Representations for Phonological Patterns	Extend Pattern to Unfamiliar Talkers (Compared to Control)	Transfer Deficit	
		Old Items	New Items
Talker-Specific	No	Yes	Yes
Talker-Specific Representations Extended to Unfamiliar Talkers	Yes	Yes	Yes
Lexicon Talker-Specific but Patterns Talker- Independent	Yes	Yes	No
Talker-Independent	Yes	No	No

Table 2

Examples of Training and Test Stimuli

Training	[netep, netepe]
	[gemit, gemite]
	[kukop, kukopo]
	[monuk, monuko]
Old Items	[gemite, *gemito]
	[*kukope, kukopo]
New Items	[bedite, *bedito]
	[mukobo, *mukobe]

Table 3

Means (Standard Deviations) for All Conditions

	Old Items			New Items		
	Familiar Talker	Unfamiliar Talker	Familiar Talker	Unfamiliar Talker	Familiar Talker	Unfamiliar Talker
Experiment 1						
Female Talker Training	0.82 (0.18)	0.70 (0.22)	0.75 (0.20)	0.68 (0.19)		
Male Talker Training	0.88 (0.16)	0.86 (0.20)	0.76 (0.14)	0.74 (0.20)		
Control	0.46 (0.12)	0.47 (0.15)	0.52 (0.13)	0.48 (0.14)		
Experiment 2						
Female Talker Training	0.84 (0.22)	0.75 (0.11)	0.64 (0.20)	0.73 (0.19)		
Male Talker Training	0.75 (0.19)	0.69 (0.24)	0.63 (0.24)	0.64 (0.20)		
Control	0.45 (0.14)	0.51 (0.16)	0.54 (0.11)	0.44 (0.12)		