The Role of Stimuli Perceptibility in Learning and Generalization of Phonological Patterns

Sara Finley

Pacific Lutheran University

finleysr@plu.edu

Author Note

Sara Finley <https://orcid.org/0000-0002-7090-8108>

We are grateful to the participants at Pacific Lutheran University and Amazon's Mechanical Turk for their cooperation and assistance, as well as to Amanda Fretland-Thomas, Paulin Schmit, Nicole Wirth, audiences at the 2018 C-SAIL conference, and anonymous reviewers for helpful comments related to this research. This work was funded in part by a Benson-Starkovich grant to SF.

I have no conflicts of interest to declare.

Correspondence concerning this article should be addressed to Sara Finley, Department of Psychology, Pacific Lutheran University, 12180 Park Ave S. Tacoma, WA 98447.

Email: finleysr@plu.edu

Abstract

Phonological patterns tend to be restrictive, often only applying to a subset of sounds in the language (e.g., height harmony applying only to front vowels). When given a limited input, what factors do learners use to determine if a novel sound participates in a phonological process? In two artificial language learning experiments, adult participants were exposed to a metathesis pattern where two stop consonants (from a six-consonant inventory) switched places with the fricative /f/ (e.g., /ep/ + /fu/ → [efpu]), and were then tested on all six consonants. Participants who were only exposed to voiced or voiceless segments showed robust generalization, but participants exposed to only velar or only labial segments showed weaker generalization. In Experiment 2, a perception task revealed a significant correlation between perceptibility of the exposure items and generalization to novel consonants. These results support theories of phonology that are informed by phonetic properties of speech sounds.

Keywords: linguistics; phonology; phonological features; artificial language learning; generalization

The Role of Stimuli Perceptibility in Learning and Generalization of Phonological Patterns

The ability to learn and generalize novel patterns is an important component to human cognition. This is particularly important for phonological rules, which rely on mappings between an underlying form and a surface form. For example, metathesis involves the switching of two elements (e.g., /p/ switches with /f/ in /ep/ + /fu/ → [efpu]). Learning a specific metathesis pattern involves encoding not only that two elements ‘switch’, but which elements do, and when. In the case of /ep/ + /fu/ → [efpu], the rule applies to the voiceless obstruents /f/ and /p/, but there is a question of how general the metathesis rule is: does it apply only to obstruents, or does it apply only to voiceless consonants, or is the pattern more general (applying to all consonants)?

Many phonological patterns (including metathesis) apply only to a particular subset of sounds. For example, consonant metathesis in Andi, a Northeast Caucasian language, applies only to /br/ → [rb] clusters (Hume, 2016). In Lummi, a Salishan language, consonant metathesis applies to all glottalized consonants in C1, and has no restrictions on C2 (Hume, 2016). Such variations on restrictiveness have led to important questions about when and how learners make generalizations to novel segments when learning phonological rules. Finley (2017ab) showed that learners are more likely to generalize to novel positions in metathesis when the novel position improves the syllable structure, in line with perceptually-based constraints (Staroverov & Finley, 2021), suggesting that learners encode the constraints that drive phonological patterns. It is unclear, however, whether learners might generalize to novel segments in a metathesis pattern when the structural motivation remains unchanged.

While some studies have shown that learners will generalize to novel segments after exposure to a single segment type (Linzen & Gallagher, 2017), other studies have shown no or very limited generalization (Cristia & Seidl, 2008), suggesting that generalization to novel

segments is constrained (Finley & Badecker, 2009). Other studies have suggested that learning and generalization is tied to perceptual factors (Kemper, 2016; White, 2014; Wilson, 2006).

The mechanisms that constrain generalization to novel segments in language learning is still unclear. Cristia et al., (2013) explored the role of similarity in generalization to novel segments. Participants were exposed to a phonotactic pattern in which a subset of obstruents from the language were displayed in onset position (e.g., /b, d, v, z/ as syllable onsets). Participants were then asked to judge the frequency of novel words. The results showed that novel segments that were highly similar to the training segments (in terms of articulatory measurements or phonological features) tended to be rated as more frequent than items that were very distinct.

These results suggest a role for similarity between training and test items as a factor of generalization, where learners generalize to items that are sufficiently similar to the items in the training set. Such a framework is consistent with domain general models of learning and generalization in psychology, where learners are more likely to generalize along a ‘gradient’ where responses to novel items are directly linked to the similarity of the exposure items (Shepard, 1987; Tenenbaum & Griffiths, 2001). Similarity has also been used as a metric for generalization in exemplar-based models of categorization (Nosofsky, 1986, 1988).

Perceptual similarity has been shown to affect learning and generalization in artificial phonological learning tasks. White (2014) showed that participants who were trained on a perceptually similar spirantization alternation (e.g., /p/ → [f]) were less likely to generalize to spirantization alternations that resulted in a larger perceptual change (e.g., /p/ → [v]), compared to participants who were trained on the perceptually distinct alternation. In addition, Wilson (2006) showed that participants were more likely to generalize to a velar palatalization pattern if

they were trained on a perceptually distinct alternation. While these studies are important for understanding similarity in learning and generalization, there are still open questions related to how learners make inferences about the generalizability of a phonological rule, based on a limited exposure set. For example, the large perceptual changes in White (2014) typically involved more than one feature, confounding perceptual and featural differences. In Wilson (2006), the pattern of generalization was not consistent across all conditions (Kapatsinski, 2008). While Cristia et al. (2013) showed a role for similarity and generalization, it was not possible to differentiate between perceptual and featural explanations. In addition, some measurements made contradictory predictions.

In order to better understand the role of sub-featural factors that play a role in learning and generalization in artificial language learning, it is important to use a relatively simple phonological pattern that can be described in terms of a single phonological feature. The present study makes use of metathesis, a phonological pattern that does not change its featural qualities from underlying form to surface form. In metathesis, two sounds switch places, as in /apfu/ → [afpu]. Metathesis has been used in previous artificial language learning experiments (Finley, 2017a, 2017b), demonstrating that American English-speaking participants can readily learn and generalize a metathesis pattern in an experimental setting.

Featural similarity is controlled in the present study because all participants are exposed to a metathesis pattern where the set of sounds can be described in terms of a single phonological natural class (e.g., voiced stops, velar stops). This also controls for the complexity of the rule, which has been known to affect both learning and generalization (Durvasula & Liter, 2020; Moreton et al., 2017), particularly since computational complexity can be confounded with featural complexity (Glewwe, 2019; Moreton & Pater, 2012; Seidl & Buckley, 2005).

In Experiment 1, adult, American English-speaking participants were exposed to a novel consonant metathesis pattern in which two stops from a six-stop inventory (/p, t, k, b, d, g/) switched places with the voiceless fricative /f/. In order to control for complexity, the two stops always formed a natural class, thus creating four training conditions: Labial, Velar, Voiceless, and Voiced. Participants were tested on all six stops, and generalization to novel stops was compared. Experiment 2 replicated Experiment 1, but also included ‘catch’ perception trials during training, where participants were asked which stop they heard (e.g., /p/ or /k/ in the Voiceless condition). Performance on the learning and generalization task was correlated with performance on the perception-based catch trials, suggesting a role for perceptual similarity of training items in learning and generalization to novel items.

Experiment 1

Method

Participants

Seventy-one participants were recruited from the participant pool at a university in the Pacific Northwest of the United States, who provided written informed consent for their participation. Participants were granted course credit for their participation. All participants were adult American English speakers, who did not participate in any previous artificial language learning study with metathesis.

Design and Materials

Participants were auditorily trained on a phonological (consonant metathesis) pattern and then tested using a two-alternative forced choice test that included segments not presented at training. The metathesis pattern was presented using the same triad design as previous experiments that set to demonstrate and input and an output form (Davidson et al., 2004; Finley,

2017a). Participants heard two single syllable items, followed by the combined form, where the two consonants had switched places (e.g., /VC₁/ + /C₂V/ → [VC₂C₁V], /ad/ + /fu/ → [afdu]). The first consonant was one of two stops, depending on the condition: Velar (/g/ and /k/), Labial (/p/ and /b/), Voiced (/b/ and /g/), and Voiceless (/p/ and /k/). The second consonant was always the voiceless fricative /f/. The training phase consisted of 12 triads repeated eight times, each in a random order.

Participants completed the two-alternative forced choice test immediately following the training phase. Each test trial consisted of two sets of /VC/ + /CV/ → [VCCV] triads, one where C₁ and C₂ switched (metathesis, /VC₁/ + /C₂V/ → [VC₂C₁V]), and one where C₁ and C₂ did not change places (no change, /VC₁/ + /C₂V/ → [VC₁C₂V]). Each response was coded as ‘correct’ if the participant chose the metathesis option. Note that both triads included identical segments, and only differed as to whether the triad indicated metathesis.

There were two types of test items: Old and New Segment. Old items included triads that were also included in the training set. New Segment items contained the four consonants not heard in training. This included /t/ and /d/ for all conditions, as well as the two consonants (from the set /p/, /k/, b/, and /d/) that was not heard in training (e.g., /p/ and /b/ in the Velar Training condition). There were 12 Old items, and 35 New Segment items. The New Segment items contained six items each with the segments: /t/, /k/, /g/, /p/, and /b/, and five with the segment /d/ (the lack of a sixth item with /d/ was due to experimenter error). All test items included the same fricative /f/ from the training phase, and the consonants were in the same location as the training phase. Examples of training and test items can be found in Table 1. Additional materials and analysis code can be found at: <https://osf.io/5s8fm/>.

Condition	Training Items	Old	New Segment
Voiced	eg fa efga	ag fo afgo vs.	ok fa ofka vs.
	ab fe afbe	ag fo *agfo	ok fa *okfa
Voiceless	ak fi afki	ok fa ofka vs.	ag fo afgo vs.
	op fu ofpu	ok fa *okfa	ag fo *agfo
Labial	ab fe afbe	eb fu efbu vs.	ag fo afgo vs.
	op fu ofpu	eb fu *ebfu	ag fo *agfo
Velar	eg fa efga	ag fo afgo vs.	eb fu efbu vs.
	ak fi afki	ag fo *agfo	eb fu *ebfu

Table 1: Examples of training and test items

Stimuli

All stimuli were recorded by an adult, female, native English speaker in a sound-attenuated booth at a 44kHz sampling rate. While the speaker was aware that the stimuli would be used for an artificial grammar learning experiment, she was unaware of the purpose of the study or the specific hypotheses that were tested. Each audio file was normalized to 70dB using Praat (Boersma & Weenink, 2017).

Procedure

The experiment was run in a traditional lab setting using PsychoScopeX (Cohen et al., 1993) (note that data for Experiment 1 was collected prior to the COVID-19 pandemic). Participants were told that they would be hearing words in sets of three: two words and their combined form, with the example from English, “tooth, brush, toothbrush”. The training was followed by the forced-choice test described above. Participants were told to choose which of two sets of three items most likely belonged to the language that they had just heard; if they believed the first set of three items most likely came from the language, they were to press the ‘a’ key, and if they believed the second set of three items most likely came from the language, they were to press the ‘l’ key. Participants were told to respond as quickly and accurately as possible, and to wait until both triads were heard before responding. The entire procedure took approximately 15 minutes.

Results

Proportions of correct (metathesis) responses can be found in Table 2, and in Figure 1. The data were fit into generalized linear mixed effect models fit by the Laplace approximation using the lme4 package in R (R Development Core Team, 2018). In order to minimize convergence errors, separate models were run for Old and New Segment items. The baseline intercept of each model was compared to 50% chance, and a significant intercept was interpreted as significantly different from chance. In order to correct for multiple comparisons, Tukey-corrected pairwise comparisons were generated using the lsmeans/emeans package (Lenth, 2016, 2018). Due to convergence errors, all models were created with random intercepts for subjects and items, unless otherwise noted.

	/p/	/t/	/k/	/b/	/d/	/g/
Labial (/pb/)	0.80	0.64	0.73	0.74	0.72	0.72
Velar (/kg/)	0.55	0.58	0.62	0.64	0.56	0.60
Voiced (/bg/)	0.69	0.71	0.75	0.80	0.68	0.88
Voiceless (/pk/)	0.77	0.74	0.87	0.74	0.73	0.78

Table 2. Mean responses to items by segment (Experiment 1). Old items are in grey.

Figure 1: Experiment 1 results (means and standard errors).

Old Items

Participants in the Voiced (mean = 0.83 ± 0.053), $\beta = 2.35$, $SE = 0.48$, $z = 3.88$, $p < 0.001$, Voiceless (mean = 0.82 ± 0.053), $\beta = 2.48$, $SE = 0.44$, $z = 3.46$, $p < 0.001$, and Labial (mean = 0.77 ± 0.057), $\beta = 1.89$, $SE = 0.48$, $z = 3.95$, $p < 0.001$) conditions all selected the correct/metathesis items for Old items at a rate significantly different from chance. However, participants in the Velar condition failed to select the correct/metathesis items for Old items at a rate greater than chance (mean = 0.60 ± 0.068), $\beta = 0.61$, $SE = 0.37$, $z = 1.63$, $p = 0.10$.

Pairwise comparisons revealed that there were significantly more correct responses to Old items in the Voiced and Voiceless conditions compared to the Velar condition. No other pairwise comparison was significant. Summaries of the Tukey corrected pairwise comparisons can be found in Table 3.

	β	SE	z	p
Velar vs. Voiced	-1.56	0.57	-2.77	0.029
Velar vs. Voiceless	-1.45	0.56	-2.59	0.048
Velar vs. Labial	-1.047	0.54	-1.93	0.21
Voiced vs. Voiceless	0.12	0.58	0.20	1.00
Voiced vs. Labial	0.52	0.57	0.91	0.80
Voiceless vs. Labial	0.40	0.56	0.71	0.89

Table 3: Tukey-corrected pairwise comparisons (Experiment 1): Old items

New Segment Items

Participants in the Voiced (mean = 0.73 ± 0.044), $\beta = 1.51$, SE = 0.53, $z = 4.69$, $p < 0.001$, Voiceless (mean = 0.75 ± 0.043), $\beta = 1.86$, SE = 0.50, $z = 3.73$, $p < 0.001$, and Labial (mean = 0.70 ± 0.044), $\beta = 1.40$, SE = 0.48, $z = 3.95$, $p < 0.001$) conditions all selected the correct/metathesis items for New Segment items at a rate significantly different from chance. Participants in the Velar condition failed to select correct responses to New Segment items at a rate greater than chance (mean = 0.58 ± 0.049), $\beta = 0.50$, SE = 0.35, $z = 1.41$, $p = 0.16$. Pairwise comparisons (summarized in Table 4) revealed no significant differences in responses to New Segment items across conditions.

	β	SE	z	p
Velar vs. Voiced	-1.04	0.60	-1.73	0.31
Velar vs. Voiceless	-1.68	0.74	-2.28	0.10
Velar vs. Labial	-0.96	0.64	-1.52	0.43
Voiced vs. Voiceless	-0.65	0.76	-0.85	0.83
Voiced vs. Labial	0.074	0.67	0.11	1.00
Voiceless vs. Labial	0.72	0.64	0.91	0.80

Table 4: Tukey-corrected pairwise comparisons (Experiment 1): New Segment items

Discussion

While there were no significant differences across conditions for New Segment items, some general trends emerged. First, participants in Experiment 1 learned and generalized a novel metathesis pattern when exposed to two consonants from a six-consonant inventory in all conditions except the Velar condition. Second, there did not seem to be any distinction between

whether the novel segment differed from the training segment by a single feature (e.g., voicing only) or by multiple features (e.g., voicing and place). This suggests that when participants generalized to novel segments, they did so freely, inferring a general rule that could apply to all stops.

While participants in the Velar condition failed to generalize to novel segments, their performance on Old items was also lower. Given the simplicity of the metathesis pattern, it is unclear why this might be. It is possible that participants in the Velar condition were less likely to attend to the pattern, either as a statistical anomaly, or due to some properties idiosyncratic to the condition. To further explore this, a second experiment was run to replicate the results of Experiment 1. This replication included ‘catch’ trials during training to determine the perceptibility of the training items from each other. For example, in the Labial condition, participants were asked whether each trial contained a ‘p’ or a ‘b’.

Experiment 2

Experiment 2 replicates Experiment 1 with a few minor changes. Participants completed the study online using the Finding Five (2019) experiment development platform, rather than in a traditional lab-based setting. Because of this, the training trials were self-paced; participants could move between items at their own pace, clicking ‘Continue’ after each trial, rather than having each item played automatically. In addition, a perception task was included in the training phase, described below. This was used both as a way to help maintain attention to the task, and to judge the perceptibility of the stops in the training set.

Method

Participants

One hundred and twenty-two adult, English-speaking participants were recruited from Amazon's Mechanical Turk, with accounts located in the United States. Participants were paid \$1.50 for their cooperation. Data from 20 participants were excluded from analysis; 16 for failure to complete the task in less than 30 minutes, and four due to experimenter error.

Design

Experiment 2 included the same design as Experiment 1, except that half of the training trials included perception-based catch trials, asking which consonant appeared in the item. For example, in the Labial condition, participants were asked whether the item contained a /p/ or a /b/, and in the Voiced condition, participants were asked whether the item contained a /b/ or a /g/.

Stimuli

The stimuli in Experiment 3 were identical to the stimuli used in Experiment 2.

Procedure

The procedure was identical to Experiment 1 except that each training item either had a perception question, or a Continue button. Thus, unlike Experiment 1, the experiment did not advance automatically. In addition, participants were given a sample sound (/udvu/) to test their headphones before beginning the experiment.

Results

Proportions of correct (metathesis) responses were recorded for each participant, and can be found in Table 5 and Figure 2. The data were analyzed in the same manner as Experiment 1, with the additional analysis of the perception trials.

	/p/	/t/	/k/	/b/	/d/	/g/
Labial (/pb/)	0.68	0.56	0.51	0.70	0.52	0.53
Velar (/kg/)	0.61	0.60	0.76	0.65	0.65	0.72
Voiced (/bg/)	0.69	0.70	0.73	0.83	0.70	0.81
Voiceless (/pk/)	0.83	0.65	0.76	0.65	0.61	0.77

Table 5. Responses to items by segment (Experiment 2). Old items are in grey.

Figure 2: Experiment 2 results (means and standard errors).

Old Items

Participants in all four conditions: the Voiced (mean = 0.81 ± 0.041), $\beta = 2.23$, SE = 0.49, $z = 5.81$, $p < 0.001$, Voiceless (mean = 0.81 ± 0.050), $\beta = 2.23$, SE = 0.46, $z = 4.88$, $p < 0.001$, Labial (mean = 0.68 ± 0.048), $\beta = 1.10$, SE = 0.29, $z = 3.83$, $p < 0.001$) and Velar (mean = 0.60 ± 0.068), $\beta = 2.11$, SE = 0.65, $z = 3.23$, $p = 0.0012$) conditions selected the metathesis pattern at a rate greater than chance for Old items. Pairwise comparisons (summarized in Table 6) revealed that there were no significant differences in correct responses to Old items across conditions.

	β	SE	z	p
Velar vs. Voiced	-0.78	0.58	-1.33	0.54
Velar vs. Voiceless	-0.72	0.64	-1.11	0.68
Velar vs. Labial	0.34	0.56	0.60	0.93
Voiced vs. Voiceless	0.062	0.62	0.10	1.00
Voiced vs. Labial	1.11	0.53	2.09	0.16
Voiceless vs. Labial	1.053	0.60	1.76	0.29

Table 6: Tukey-corrected pairwise comparisons (Experiment 2): Old items

New Segment Items

Participants in the Voiced (mean = 0.69 ± 0.035), $\beta = 1.33$, SE = 0.36, $z = 3.72$, $p < 0.001$ and Voiceless (mean = 0.68 ± 0.042), $\beta = 1.38$, SE = 0.43, $z = 3.213$, $p = 0.0013$, conditions were conditions all selected the correct/metathesis items for New Segment items at a rate significantly different from chance. Participants in the Velar (mean = 0.58 ± 0.049), $\beta = 1.65$, SE = 0.62, $z = 1.87$, $p = 0.062$, and Labial (mean = 0.52 ± 0.037), $\beta = 0.20$, SE = 0.27, $z = 0.73$, $p = 0.47$, conditions failed to select the correct item at a rate significantly greater than chance. Pairwise comparisons (summarized in Table 7) revealed that there were no significant differences in correct responses to New Segment items across conditions.

	β	SE	z	p
Velar vs. Voiced	-0.30	0.56	-0.53	0.95
Velar vs. Voiceless	-0.26	0.63	-0.43	0.97
Velar vs. Labial	0.94	0.56	1.69	0.33
Voiced vs. Voiceless	0.031	0.59	0.053	1.00
Voiced vs. Labial	1.24	0.51	2.44	0.069
Voiceless vs. Labial	1.21	0.58	2.09	0.16

Table 7: Tukey-corrected pairwise comparisons (Experiment 2): New Segment items

Perception Trials

Because the perception trials were a generally easy task (e.g., no noise), it was expected that participants would perform around ceiling on this task. All responses showed a mean above 80%: Velar (mean = 0.83 ± 0.23), Labial (mean = 0.86 ± 0.019), Voiced (mean = 0.94 ± 0.014), and Voiceless (mean = 0.95 ± 0.014). Pairwise comparisons (summarized in Table 8) revealed

that the Voiced and Voiceless conditions showed significantly higher rates of correct responses to both Velar and Labial conditions, but there were no differences between Voiced and Voiceless conditions, or Velar and Labial conditions.

	β	SE	z	p
Velar vs. Voiced	-3.31	0.64	-5.13	<0.001
Velar vs. Voiceless	-2.27	0.77	-2.95	0.017
Velar vs. Labial	-0.20	0.59	-0.34	0.99
Voiced vs. Voiceless	1.03	0.77	1.35	0.53
Voiced vs. Labial	3.11	0.62	5.06	<0.001
Voiceless vs. Labial	2.08	0.73	2.84	0.023

Table 8: Tukey-corrected pairwise comparisons (Experiment 2): Catch trials

In addition, a Pearson correlation test showed a significant correlation between correct responses to the perception items and correct responses to New Segment items (across all conditions, combined), $r = 0.38$, $p < 0.001$, as shown in Figure 3.

Figure 3: Correlation between perception task and correct responses to New Segment items

The results of the correlation demonstrate that participants who scored well on the perception catch trials were more likely to generalize to novel segments. While the overall scores were generally very high, the differences may reflect some participants' ability to properly identify the sounds heard in training.

Discussion

Experiment 2 generally replicated Experiment 1 with some minor differences; participants in the Velar condition were above chance for the Old Segment items, and participants in the Labial condition were not above chance for New Segment items. The results of perception trials were in line with the generalization results; there was a significant correlation between generalization and performance on the perception task, and there were significantly higher rates of correct responses on the Catch trials for the conditions that generalized to novel segments (Voiced and Voiceless) compared to the conditions that did not (Velar and Labial).

There are two possible explanations for the correlation between the catch trials and the responses to New Segment items. One is that participants who struggled to perceive the differences between the sounds at training also struggled to learn the rule, and hence, generalize to novel segments. A second possibility is that participants who were not attending to the task (e.g., just ‘clicking’ through) were also less likely to learn the pattern, or attend to the test item as well. If the correlation between performance on the perception task and performance on the two alternative forced choice test is just based on attentional factors, there is a question of why the responses were lower in the Labial and Velar conditions, particularly since participants in the Velar condition failed to generalize to novel segments in Experiment 1.

General Discussion

The present study presented two artificial grammar learning experiments testing for generalization to novel segments in phonological pattern learning. Adult, English-speaking participants were exposed to a consonant metathesis pattern in which a stop consonant switched places with a voiceless fricative (e.g., /ep/ + /fu/ → [efpu]). Participants were exposed to two consonants from a six-vowel inventory, and were tested on their learning and generalization to

novel segments. Participants who were exposed to either all Voiced or Voiceless sounds generalized to novel segments in both experiments, but participants in the Velar (Experiments 1 and 2) and Labial (Experiment 2) did not. This may be due to the perceptibility of the consonants in the training phase. In Experiment 2, participants were given ‘catch’ perception trials where they had to choose which consonant was in each trial. Generalization to novel segments was significantly correlated with performance on this task, and participants in the Voice conditions showed significantly higher rates of correct responses compared to participants in the Place conditions, in line with successful generalization.

Assuming that the lower levels of perceptibility in the Labial and Velar conditions in Experiment 2 was not due to random chance (e.g., more participants in these conditions just happened to fail to attend to the task), there is still a question of *why* the stimuli in these conditions should be less perceptible. Because the overall set of items are identical between a combined Voiced and Voiceless conditions (e.g., /p, k, b, g/) and a combined Labial and Velar conditions (also /p, k, b, g/), the failure to generalize to Velar and Labial conditions in Experiment 2 is unlikely to be driven by lower perceptibility of some items compared to others. The difference between each condition was which phonemes the participant was asked to differentiate between (e.g., /p/ and /b/ in the Labial condition and /p/ and /k/ in the Voiceless condition). Previous research has shown that voicing contrasts (e.g., /p/ vs. /b/ are more perceptible in noise compared to place contrasts (e.g., /p/ vs. /k/) (Benkí, 2003; Miller & Nicely, 1955; Wang & Bilger, 1973). These findings cannot explain the present results, since voicing contrasts in the Labial and Velar condition were actually less perceptible than the place contrasts in the Voiced and Voiceless conditions. However, it is important to note that in these previous studies, participants were tested on CVC syllables, not stop-fricative clusters. In the present

study, coarticulation from the /f/ in the /f/+stop clusters could have impaired the perceptibility of the voicing contrasts, which would have impaired perception of items in the Labial and Velar conditions, whose items only differentiated in terms of voicing. If the voicing contrast was obscured in the study participants in the Voiced condition may have perceived some items to be voiceless (and vice-versa for the voiceless condition), which would have made generalization from voiced to voiceless (and vice versa) easier because some participants would not have perceived the segments to be novel. This is consistent with previous results that showed that participants generalized metathesis based on perceptual constraints, particularly in avoiding voiced obstruent codas (Finley, 2017b).

Another possible explanation for the lower perceptibility scores in the Labial and Velar conditions is that both the Voiced and Voiceless conditions made use of labial and velar sounds, which, from an articulatory point of view, take up greater space within the vocal tract. Since alveolar segments are intermediate in the vocal tract space, participants may have seen generalization to alveolars as interpolation, rather than extrapolation, which may be easier for learners to accomplish (DeLosh et al., 1997).

In both Experiments 1 and 2, the Velar and the Labial conditions had numerically overall lower responses to New Segment items compared to the Voiced and Voiceless conditions. However, the individual comparisons failed to reach statistical significance with corrections for multiple comparisons. It may be possible to increase statistical power by combining the Voiced and Voiceless conditions into a single ‘Voice’ condition and the Velar and Labial conditions into a ‘Place’ condition. Such collapsing of conditions might be justified given the overall differences in perceptibility of place of articulation and voicing. If participants use different strategies for perceiving place of articulation than voicing, then it may make sense that similar patterns of

results will be found in the Place and Voice conditions. To test whether Place and Voice showed different levels of generalization to novel segments, separate models with random intercepts for subjects and items, and random slopes for condition by subject tested the degree to which generalization occurred in the Place and Voice conditions. For Experiment 1, the difference was only marginally significant, $\beta = 0.86$, $SE = 0.50$, $z = 1.74$, $p = 0.083$, and for Experiment 2, the difference was significant, $\beta = 0.85$, $SE = 0.39$, $z = 2.29$, $p = 0.029$. These results support the possibility that participants are more likely to generalize to novel segments if they are trained on items that differ in terms of place of articulation (but have the same voicing).

These results have important consequences for theories of phonological pattern learning. First, the present study provides an example of a replication of a phonological pattern learning from a traditional laboratory setting (Exp 1) to an online setting (Exp 2). While there were some differences in the results, the general patterns remained intact. In addition, the high proportion of correct responses to the catch perception trials demonstrates that participants could hear and attend to the auditory stimuli. While there are studies that made use of online learning for phonological patterns (Linzen & Gallagher, 2017; Martin & White, 2021; Staroverov & Finley, 2021), as well as studies that have replicated lab experiments (Ramsey et al., 2016), few studies have replicated phonological pattern learning in both a laboratory and an online setting. These results should provide confidence in using online crowdsourcing for artificial language learning experiments in phonology, which has become increasingly popular due to the COVID-19 pandemic.

Second, the present results add to the growing body of literature that suggests that learners of both phonotactic patterns and phonological alternations can successfully generalize to novel segments after a short period of training, but that participants do not always generalize to

novel segments (Cristia & Seidl, 2008; Finley, 2011; Finley & Badecker, 2009; Wilson, 2006).

While complexity and featural similarity between the trained items and the novel items has been shown to play a role in generalization (Cristia et al., 2013; White, 2014), they are unlikely to explain the present results. First, all conditions could be described using a single natural class. Second, in the Voiced and Voiceless conditions, participants generalized outside of the natural class they were trained on. In addition, participants in the Velar condition in Experiment 1 showed no differences between Old and New Segment items.

In addition, the present studies provide important contributions to understanding learning and generalization more broadly, particularly how learners are able to make broad generalization with limited input (Goldberg et al., 2004; Harmon & Kapatsinski, 2017; Wonnacott, 2011). This is important because language learners are rarely given the entire set of possibilities for linguistic patterns, but mechanisms that allow the learner to make a strong inference given a limited input allows for rich, but constrained generalization (Xu & Tenenbaum, 2007).

Previous research on category learning and generalization has shown that diversity of a category can lead to both better learning and generalization both linguistically (Reeder et al., 2013) and non-linguistically (Hahn et al., 2005). While the present study did not manipulate the diversity of the category, the similarity results suggest that when the segments that undergo a phonological rule are similar, participants may be less likely to extract regularities from the rule, increasing the difficulty of learning and generalization.

In addition, there is a question of the robustness of perceptual similarity in phonological rule learning. For example, White (2014) and Wilson (2006) made use of more general studies of perceptual similarity in order to determine the predictions for generalization, but there is a possibility that perceptibility of the specific training stimuli may affect learning and

generalization (Gafos, 2021). In addition, articulatory factors that were not considered in previous studies could also play a role in learning and generalization (Cristia et al., 2013; Do & Havenhill, 2021; Kapatsinski, 2008, 2010, 2013).

While this study has important implications for both phonological theory and language learning more generally, there are several limitations that should be noted. First, while the sample size used in the present study is consistent with previous literature, it may be too small to detect fine-grained differences across conditions, with more conservative statistical models. The small sample size makes it difficult to make firm conclusions because some comparisons were not statistically significant with corrections for multiple comparisons, and it is also not clear if some trends (e.g., generally lower rates of correct responses in the Labial and Velar conditions) were due to random noise, or substantive differences in these conditions. Future research on this type of work should include a much larger sample size, particularly if data is to be collected online, where participant engagement is more difficult to assess and control.

A second limitation to the present study is that the use of English-speaking participants may have biased learners in their responses to the metathesis pattern. Since metathesis is common in several English dialects, participants are likely to be aware of metathesis prior to the study's start, which may have biased their learning and generalization of the pattern in unexpected ways. However, the rate of correct responses, particularly in the Velar and Labial conditions, was somewhat low for a relatively simple pattern. This suggests that participants did have to learn a novel pattern. In addition, while metathesis is found in English, the specific linguistic environment for metathesis (e.g., /stop/ + /f/) is not found in English, meaning that learners had to form a novel metathesis pattern, even if they had some familiarity with metathesis

prior to the experiment. While the study has its limitations, this work does present an important contribution to understanding learning and generalization of phonological patterns.

Conclusion

The present study presents two artificial language learning studies that show that learners readily generalize to novel segments after training on a novel phonological alternation, but that this generalization is constrained. One possible mechanism for this constraint on generalization is the nature of the features in the training set, particularly the confusability of the sounds heard in training; more confusable training items may lead to both degraded learning, and lower rates of generalization.

References

- Benkí, J. R. (2003). Analysis of English nonsense syllable. *Phonetica*, 60, 129–157.
<https://doi.org/10.1159/000071450>
- Boersma, P., & Weenink, D. (2017). *Praat: Doing phonetics by computer*.
- 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. <https://doi.org/10.3758/BF03204507>
- Cristia, A., Mielke, J., Daland, R., & Peperkamp, S. (2013). Similarity in the generalization of implicitly learned sound patterns. *Laboratory Phonology*, 4(2). <https://doi.org/10.1515/lp-2013-0010>
- Cristia, A., & Seidl, A. (2008). Is infants' learning of sound patterns constrained by phonological features? *Language Learning and Development*, 4(3), 203–227.
- Davidson, L., Smolensky, P., & Jusczyk, P. (2004). The initial and final states: Theoretical

- implications and experimental explorations of richness of the base. In R. Kager, J. Pater, & W. Zonnefeld (Eds.), *Constraints in phonological acquisition* (pp. 321–368). Cambridge University Press.
- DeLosh, E. L., Busemeyer, J. R., & McDaniel, M. A. (1997). Extrapolation: The sine qua non for abstraction in function learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(4), 968–986. <https://doi.org/10.1037//0278-7393.23.4.968>
- Do, Y., & Havenhill, J. (2021). Production and substantive bias in phonological learning. *Proceedings of the Annual Meetings on Phonology*, 9, 1–12. <https://doi.org/10.3765/amp.v9i0.4925>
- Durvasula, K., & Liter, A. (2020). There is a simplicity bias when generalizing from ambiguous data. *Phonology*, 850(April), 1–64.
- FindingFive Team. (2019). *FindingFive: A web platform for creating, running, and managing your studies in one place*. FindingFive Corporation (nonprofit),. <https://www.findingfive.com>
- Finley, S. (2011). Generalization to novel consonants in artificial grammar learning. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (pp. 3268–3273). Cognitive Science Society.
- Finley, S. (2017a). Learning metathesis: Evidence for syllable structure constraints. *Journal of Memory and Language*, 92. <https://doi.org/10.1016/j.jml.2016.06.005>
- Finley, S. (2017b). Perceptually Based Constraints and Metathesis: Evidence from Artificial Grammar. *Proceedings of the Annual Meetings on Phonology*, 4, 1–10. <https://doi.org/10.3765/amp.v4i0.3985>
- Finley, S., & Badecker, W. (2009). Artificial language learning and feature-based generalization. *Journal of Memory and Language*, 61(3), 423–437.

<https://doi.org/https://doi.org/10.1016/j.jml.2009.05.002>

- Gafos, A. I. (2021). Information flow, artificial phonology and typology. *Proceedings of the Society for Computation in Linguistics (SCiL)*, 4(148–157).
<https://doi.org/doi.org/10.7275/6zx1-p517>
- Glewwe, E. (2019). *Bias in phonotactic learning: Experimental studies of phonotactic implicational*. UCLA.
- Goldberg, A. E., Casenhiser, D. M., & Sethuraman, N. (2004). Learning argument structure generalizations. *Cognitive Linguistics*, 15(3), 289–316.
<https://doi.org/10.1515/cogl.2004.011>
- Hahn, U., Bailey, T. M., & Elvin, L. B. C. (2005). Effects of category diversity on learning, memory, and generalization. *Memory and Cognition*, 33(2), 289–302.
<https://doi.org/10.3758/BF03195318>
- Harmon, Z., & Kapatsinski, V. (2017). Putting old tools to novel uses: The role of form accessibility in semantic extension. *Cognitive Psychology*, 98, 22–44.
<https://doi.org/10.1016/j.cogpsych.2017.08.002>
- Hume, E. (2016). *Metathesis in Language*. <http://metathesisinlanguage.osu.edu/>
- Kapatsinski, V. (2008). Rule reliability and productivity: Velar palatalization in Russian and artificial grammar. *Laboratory Phonology*, 11.
- Kapatsinski, V. (2010). Velar palatalization in Russian and artificial grammar: Constraints on models of morphophonology. *Laboratory Phonology*, 1(2), 361–393.
- Kapatsinski, V. (2013). Conspiring to mean: Experimental and computational evidence for a usage-based harmonic approach to morphophonology. *Language*, 89(1), 110–148.
<https://doi.org/10.1353/lan.2013.0003>

- Kimper, W. A. (2016). Asymmetric generalisation of harmony triggers. In G. Ó. Hansson, A. Farris-Trimble, K. McMullin, & D. Pulleyblank (Eds.), *Proceedings of the 3rd Annual Meeting on Phonology* (pp. 1–10). Linguistic Society of America.
- Lenth, R. (2016). *Using lsmeans*. <https://cran.r-project.org/web/packages/lsmeans/vignettes/using-lsmeans.pdf>
- Lenth, R. V. (2018). *emmeans: Estimated Marginal Means, aka Least-Squares Means*.
- Linzen, T., & Gallagher, G. (2017). Rapid generalization in phonotactic learning. *Laboratory Phonology: Journal of the Association for Laboratory Phonology*, 8(1). <https://doi.org/10.5334/labphon.44>
- Martin, A., & White, J. C. (2021). Vowel harmony and disharmony are not equivalent in learning. *Linguistic Inquiry*, 52(1), 227–239. https://doi.org/10.1162/ling_a_00375
- Miller, G., & Nicely, P. (1955). An analysis of perceptual confusions among some English consonant. *Journal of the Acoustical Society of America*, 27(2), 338–352. <https://doi.org/10.1121/1.1907526>
- Moreton, E., & Pater, J. (2012). Structure and substance in artificial-phonology learning, part I: Structure. *Language and Linguistics Compass*, 6(11), 686–701. <https://doi.org/https://doi.org/10.1002/lnc3.363>
- Moreton, E., Pater, J., & Pertsova, K. (2017). Phonological concept learning. *Cognitive Science*, 41(1), 4–69. <https://doi.org/10.1111/cogs.12319>
- Nosofsky, R. M. (1986). Attention, Similarity, and the Identification-Categorization Relationship. *Journal of Experimental Psychology: General*, 115(1), 39–57. <https://doi.org/10.1037/0096-3445.115.1.39>
- Nosofsky, R. M. (1988). Similarity, Frequency, and Category Representations. *Journal of*

Experimental Psychology: Learning, Memory, and Cognition, 14(1), 54–65.
<https://doi.org/10.1037/0278-7393.14.1.54>

R Development Core Team, R. (2018). R: A language and environment for statistical computing.
In R. D. C. Team (Ed.), *R Foundation for Statistical Computing* (Vol. 1, Issue 2.11.1, p. 409).
R Foundation for Statistical Computing. <https://doi.org/10.1007/978-3-540-74686-7>

Ramsey, S. R., Thompson, K. L., McKenzie, M., & Rosenbaum, A. (2016). Psychological research
in the internet age: The quality of web-based data. *Computers in Human Behavior*, 58, 354–
360. <https://doi.org/10.1016/j.chb.2015.12.049>

Reeder, P., Newport, E. L., & Aslin, R. (2013). From shared context to syntactic categories: The
role of distributional information in learning linguistic form-classes. *Cognitive Psychology*,
66, 30–54.

Seidl, A., & Buckley, E. (2005). On the learning of arbitrary phonological rules. *Language
Learning and Development*, 1, 289–316.

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science.
Science, 237(4820), 1317–1323.

Staroverov, P., & Finley, S. (2021). Learning Syllable-Based vs. Segment-Based Infixation
Patterns. *WCCFL*, 38, 397–405.

Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference.
The Behavioral and Brain Sciences, 24, 629–640; discussion 652–791.
<https://doi.org/10.1017/S0140525X01000061>

Wang, M. D., & Bilger, R. C. (1973). Consonant confusions in noise: A study of perceptual
features. *The Journal of the Acoustical Society of America*, 54(5), 1248–1266.

White, J. C. (2014). Evidence for a learning bias against saltatory phonological alternations.

Cognition, 130(1), 96–115. <https://doi.org/10.1016/j.cognition.2013.09.008>

Wilson, C. (2006). Learning phonology with substantive bias: An experimental and computational study of velar palatalization. *Cognitive Science*, 30, 945–982. https://doi.org/10.1207/s15516709cog0000_89

Wonnacott, E. (2011). Balancing generalization and lexical conservatism: An artificial language study with child learners. *Journal of Memory and Language*, 65(1), 1–14. <https://doi.org/10.1016/j.jml.2011.03.001>

Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, 114(2), 245–272. <https://doi.org/10.1037/0033-295X.114.2.245>