

Unexpectedness makes a sociolinguistic variant easier to learn: An alien-language-learning experiment

Wei Lai (weilai@sas.upenn.edu)

Department of Linguistics, University of Pennsylvania

Péter Rác (RaczP@ceu.edu)

Cognitive Development Center, Central European University

Gareth Roberts (gareth.roberts@ling.upenn.edu)

Department of Linguistics, University of Pennsylvania

Abstract

We report two artificial-language-learning experiments investigating if the acquisition of sociolinguistic associations is facilitated by two kinds of expectation violation: encountering a variant (a) for the first time or (b) in an ungrammatical context. Participants learned an artificial language with two dialects, each spoken by one of two alien species: *Gulus* and *Norls*. The two dialects differed with regard to a plural suffix: *Gulus* mostly used *-dup*, and *Norls* mostly used *-nup*. In the first learning phase, participants learned the language without aliens; in the second learning phase, they were exposed to it with alien interlocutors. In Experiment 1 we manipulated whether *-nup* occurred in the first learning phase; in Experiment 2 we manipulated linguistic constraints on its occurrence. The acquisition of sociolinguistic association was evaluated by asking participants to select suffixes given aliens and vice versa. We found that sociolinguistic acquisition was facilitated in Experiment 1, but not Experiment 2. In Experiment 2, however, a post hoc analysis revealed that participants who had learned the grammatical context of the linguistic conditioning did experience facilitation, while those who had not did not. Our results provide laboratory evidence that unexpectedness facilitates the learning of sociolinguistic variation.

Keywords: artificial-language learning; social meaning; sociolinguistics; salience; surprisal

Introduction

The role of *salience* in the acquisition and propagation of linguistic variants has long been documented in classic sociolinguistic research (Labov, 1972). Variants with higher salience are encoded with more attention and higher meta-linguistic awareness, leading them to be more easily recognized and retained than other variants with equal frequency, resulting in an acquisition bias that cannot be explained by frequency of exposure alone (Jaeger & Weatherholtz, 2016).

In this study we investigated the role of salience in facilitating the learning of sociolinguistic meaning (i.e., the association of particular linguistic variants with particular social groups). We focused in particular on the effect of previous experience on salience. Previous work has paid much attention to the role of certain kinds of *non-linguistic* experience such as social and developmental experience (Foulkes & Docherty, 2006) or social stereotypes (Levy, 2008), but *linguistic* experience is relatively understudied.

In particular, there is very little work on how the perceived salience of a sociolinguistic variant is affected by prior experience of that variant in other contexts. Jaeger and Weatherholtz (2016, p. 1) proposed that salience related to language experience can be understood in terms of expectation violation, analogous to the well-attested novelty bias effect: Novel items and events that we do not expect tend to stand out. Jaeger and Weatherholtz (2016) argued that this might occur for linguistic variants that a listener has not encountered before and might thus lead to surprisal. The salience generated by surprisal may facilitate learning the variant and its socioindexical meaning.

Although Jaeger and Weatherholtz's (2016) approach to experience-based salience seems appealing for its operationalization of expectation-related salience in an information-theoretic framework (Shannon, 1948; Hale, 2001; Levy, 2008), it is not yet supported by linguistic data. Several experimental studies on language processing show that less expected words and structures take longer to process and at greater cost (McRae, Spivey-Knowlton, & Tanenhaus, 1998; McDonald & Shillcock, 2003), with similar effects observed in comprehension tasks (Kaschak & Glenberg, 2004; Squires, 2014; Fraundorf & Jaeger, 2016). However, additional processing for novel variants in comprehension does not necessarily result in better performance in noticing or memorizing these variants or in associating them with the right social group.

The present study

The present study investigates the hypothesis that experience-dependent salience can arise from expectation violation, and cause a sociolinguistic variant to be more learnable. We used an "alien language" learning paradigm in which participants first learned a miniature artificial language and were then exposed to it in a simple social context with "alien interlocutors". We investigated two kinds of expectation violation, hypothesizing that participants would be more likely to learn a sociolinguistic association if (Experiment 1) they had not encountered it before and (Experiment 2) they had encountered it before, but subject to grammatical constraints that now appear to be violated.

As an example of the first kind of violation, one might imagine an American English speaker visiting Liverpool and hearing, for the first time, *book* pronounced with a final velar fricative [x] (as in German *Buch*) instead of the expected velar stop [k]. As an example of the second kind of violation, consider a speaker who has heard *-th* pronounced as [f], but only at the end of syllables (as in [boʊf] for *both* or [ˈɛfnɪk] for *ethnic*). For a speaker who had acquired syllable-finality as a constraint on this variant, hearing [fɪnk] for *think* would likely be relatively salient.

In both our experiments a certain variant of the alien language was associated predominately (but not necessarily exclusively) with a particular species of alien. We evaluated whether participants had learned this sociolinguistic association pattern by asking them at the end of the experiment to (a) select the variant that a given alien would most likely produce, and (b) select the alien most likely to produce a given variant. We predicted that increased salience, via expectation violation, would lead participants to do better at both tasks (though, because we did not make the sociolinguistic relationship categorical, we did not expect that the response to the two tasks would be identical).

Finally, we predicted that listeners who learned a sociolinguistic association in the experiment could generalize that relationship to new words.

Experiment 1: First encounter

Experiment Overview

In Experiment 1, we investigated whether encountering a linguistic variant for the first time in a social context would facilitate the sociolinguistic learning of that variant (we will term this hypothesized effect “first-encounter facilitation”). Participants were trained on an alien language with two dialects, each used by a different alien species, the *Gulus* or the *Norls* (Fig. 1). The dialects differed with regard to a plural suffix: *Gulus* used *-dup* as the only form of the plural suffix whereas *Norls* sometimes used *-dup* but mostly used *-nup*.

The experimental procedure consisted of three phases: Participants were first trained on the language without seeing any aliens, which was intended to establish *prior experience* with the language; then (having been introduced to the two alien species) they were exposed further to the language with alien interlocutors, which allowed them to learn associations between plural suffixes and alien species. In the third and final phase, acquisition of sociolinguistic variants was evaluated on the basis of whether participants could infer which alien might have used a given suffix and, conversely, which suffix a given alien might have used.

Crucially, we manipulated participants’ *prior experience* with the variant *-nup* such that half the participants would never be exposed to it in the first phase, encountering only *-dup* (*NoExposure* condition), whereas the other half would see both suffixes in every phase (*Exposure* condition). We predicted that participants with no experience of *-nup* in the first learning phase would find it more salient in the second

phase and better learn to associate it with *Norls*.

Method

Participants 100 participants completed Experiment 1 online within the specified amount of time (1.5 hours). After excluding participants whose duration was below the 2.5% quantile or above the 97.5% quantile of all participants, we used the data of the remaining 93 participants. There were 51 female and 43 male participants, aged 17–73 (mean: 28.9) years. 30 of them were recruited from the University of Pennsylvania subject pool (in return for course credit) and 64 were recruited through the Prolific Academic website (and were paid \$5 each). 49 participants were in the *Exposure* condition and 45 in the *NoExposure* condition.

Alien language The artificial language was composed of 14 word stems, as shown in Table 1, and a plural suffix with two variants, *-dup*, *-nup*.

Table 1: 14 Stem Words in the Alien Language

| |
|---|
| <i>nesel, laniz, firot, hiwen, maqub, jemulok, geguzis, tugan, nuwik, falon, wumos, wukin, sehilod, takoles</i> |
|---|

The 14 stem words were randomly generated by combining one or two CV syllables with a word-final CVC syllable from a segment pool of five vowels /a e i o u/ and 12 consonants /k g q h m n t s z j l w/.

Aliens The language was used by two alien species: *Gulus* and *Norls*. The stem forms were the same across dialects, but the suffix variants had different distributions: *Gulus* attached *-dup* to all 14 words to signal plurality, whereas *Norls* attached *-dup* to only four of the words (*hiwen*, *wukin*, *jemulok* and *wumos*) and *-nup* to the remaining eight words. Put differently, *Gulus* used the *-dup* variant 100% of the time as their plural suffix whereas *Norls* used *-dup* and *-nup* at a ratio of 71% to 29%. Within each alien species, six idiosyncratic aliens were designed in order to ensure that the linguistic variation on the group level wouldn’t be mistaken for variation on the individual level (See Fig. 1 for examples).



Figure 1: Alien Species: *Gulus* (left) and *Norls* (right)

Procedure The experimental task was composed of two learning phases, in which participants were trained on the alien words through passive exposure to word-object¹ pairs and multiple-choice exercises with feedback, and a test

¹We thank Professor Janet Pierrehumbert for making images of the objects available for use. The artworks are copyrighted to Northwestern University and used with permission.

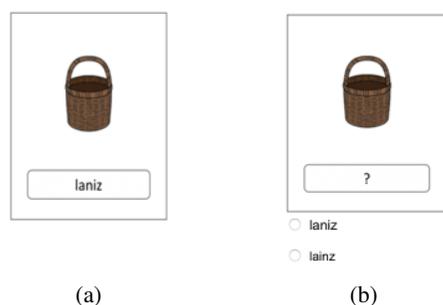


Figure 2: Example trials in learning phase 1: (a) a passive exposure trial; (b) a forced-choice trial

phase that evaluated how well participants had learned the association between plural variants and alien species.

First learning phase: Learning without aliens. The experiment started with a learning phase that exposed participants to the words of the language without any aliens. This was designed to give participants exposure to the language before introducing it in a social context. It consisted of a series of trials, with two kinds of trial, as shown in Fig. 2. In passive exposure trials (Fig. 2a) a word was paired with an image of the object(s) it referred to. Participants were instructed to memorize the word and its meaning before proceeding to the next trial. In forced-choice trials (Fig. 2b) participants had to choose the correct word to go with an image; there were always two options to choose from, one correct and one a foil generated by changing one or two segments in the stem of the correct form. Participants received one point for each correct response and no point for an incorrect one (maximum: 168 points). Feedback on the correct form and the point received for each question was provided immediately afterwards.

Participants were trained on 28 alien words (14 singular, 14 plural), which were divided into seven sets of four words each, with the constraint that each set contained two singular words and two plural words that all had different stems. For each set, a participant would see a passive exposure trial for each word in turn; then they would see a forced-choice trial for each of the same four words. Then they would proceed to the next set. The order of the seven word sets was randomized, as was the order of the four trials within each passive exposure section and within each forced-choice section. The whole process was repeated once participants had completed training on all seven word sets. In total, participants were exposed to 14 words \times 2 forms (singular and plural) \times 2 trial types (exposure, forced choice) \times 2 repetitions = 112 trials.

Alien introduction. After the first learning phase, the aliens were introduced. Participants were first presented with images of Gulus and Norls, each labeled with the species name; then the labels were removed and participants were instructed to drag and drop each alien into one of the two boxes labeled *Gulu* and *Norl*. Feedback was provided after

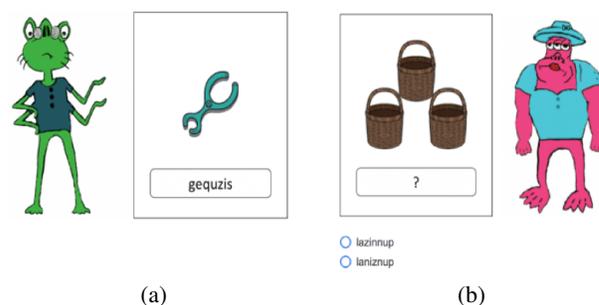


Figure 3: Example trials in learning phase 2: (a) a passive exposure trial; (b) a forced-choice trial

the drag-and-drop.

Second learning phase: Learning with aliens. After the aliens had been introduced, the second learning phase started. This phase resembled the first phase in its structure, except that each trial (whether passive-exposure or forced-choice trial) included a picture of an *alien interlocutor*, as shown in Fig. 3. Participants saw both the Gulu and the Norl form of every word, so the second learning phase was twice as long as the first learning phase (with each set of trials containing eight words rather than four). In total each participant was exposed to 14 words \times 2 forms (singular, plural) \times 2 species (Gulu, Norl) \times 2 trial types (exposure, forced choice) \times 2 repetitions = 224 trials.

Test phase: Measuring acquisition. After the second learning phase, the test phase began, which evaluated the extent to which participants had established associations between alien groups and plural suffixes. The test phase contained two tasks: a suffix-identification task in which participants had to choose which form might be used based on the presented alien interlocutor, and an alien-identification task in which participants had to choose which alien was most likely to have said a prompt word. Trials in these tasks contained both *old word* stimuli from the learning phase and *new word* stimuli that participants had never seen, in order to evaluate the generalization of sociolinguistic associations to novel items. Trial order was randomized for each participant, and the order of the two options within each trial was counterbalanced. No feedback was provided.

In *suffix identification*, trials on old words worked like forced-choice trials in the second learning phase (Fig. 3b), except that the optional answers had identical stems and different suffixes (i.e., the reverse of the situation in the learning phases). Participants were instructed to choose the form the pictured alien would likely use. Trials on new words were different: Participants were presented with a singular word, an image of the object it referred to, and an alien interlocutor; they were required to choose between a *dup*-ending word and a *nup*-ending word as the plural form (Fig. 4).

In all, the task included 56 trials on old words (14 words \times 2 species \times 2 repetitions = 56 trials), 24 trials on

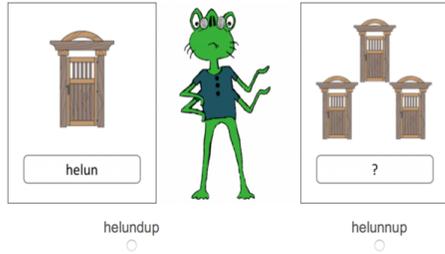


Figure 4: Example suffix-identification trial with a new word

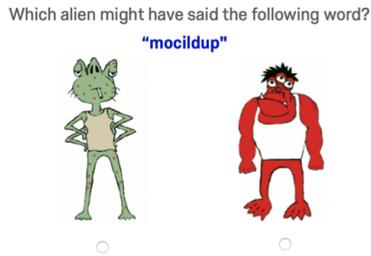


Figure 5: Example alien-identification trial

new words (6 new words \times 2 species \times 2 repetitions = 24 trials), and 34 filler trials, which tested participants on either singular words or plural words with incorrect stems.

In *alien identification* (Fig. 5) participants were given a plural word and had to choose between a Gulu and a Norl as the likely speaker of the word. The idiosyncratic aliens were kept consistent throughout the whole task, but whether they appeared on the left or the right was counterbalanced across questions. The stimulus words were generated by affixing the 14 old words and the six new words once each with *-dup* and once each with *-nup*, so that there were 40 trials (14 old words \times 2 suffixes + 6 new words \times 2 suffixes) in total.

Experimental conditions. Participants were randomly assigned to two experimental conditions: the *NoExposure* condition and the *Exposure* condition. Fig. 6 shows the distribution of variants in the two learning phases in different experimental conditions.

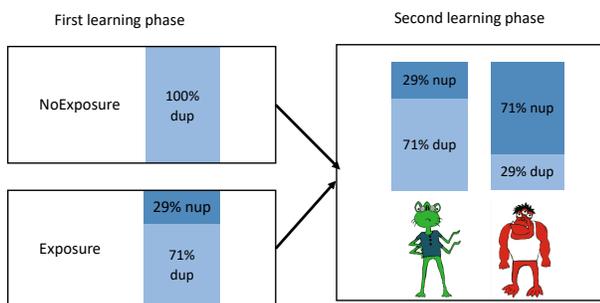


Figure 6: Variant distribution in the learning phases of Experiment 1

The two conditions differed with respect to the presence or absence of the variant *-nup* in the first learning phase: For participants in the *NoExposure* condition, plural words in this phase would always be affixed with *-dup*, whereas participants in the *Exposure* condition would see ten instances of plurals with *-dup* (71%) and four with *-nup* (29%). The two conditions were identical in the second learning phase: Gulu exclusively used *-dup* while Norls used *-nup* 71% of the time and *-dup* 29% of the time.

Results

Analyses were conducted using the R Statistical environment (R Core Team, 2014); linear models were run using the lme4 library (Bates, Mächler, Bolker, & Walker, 2015), and plots were created using ggplot (Wickham, 2016).

On average, it took participants (outliers excluded) 52 minutes ($sd = 14$) to complete the whole experiment. Out of a maximum of 168 points, participants achieved an average score of 153 ($sd = 13$).

Fig. 7 shows the aggregate results for suffix identification (left) and alien identification (right). The left panel shows how often participants selected the *-nup* suffix for a given alien the suffix-identification task. Consistent with our predictions, participants in the *NoExposure* group were more inclined to choose a *-nup* word for a Norl than those in the *Exposure* condition. Notably, the *-nup* response ratios given a Norl were relatively low in both conditions, nowhere matching the 71% in the input. The right panel shows what proportion of the time participants selected a Norl for a given suffix in the alien-identification task. Again, consistent with the hypothesis, participants in the *NoExposure* condition were more inclined to choose a Norl given a *-nup* word and to choose a Gulu given a *-dup* word, compared with those in the *Exposure* condition, who chose Norl interlocutors for both *-dup* and *-nup* at chance level.

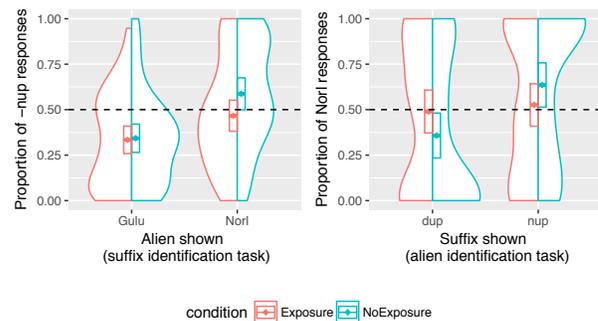


Figure 7: Proportion of *-nup* responses in suffix identification and Norl responses in alien identification (including 95% confidence interval). Dotted line indicates chance level.

Mixed-effects logistic regression models were fit on the two tasks, with Response as the dependent variable, Condition (Exposure as the intercept), Stimulus (Norl as the intercept in suffix identification; *-dup* as the intercept

in alien identification) and their interactions as independent variables, and Participant and Word as random factors. Both models revealed a significant Condition effect ($\beta = 0.57, p = 0.001$ in suffix identification; $\beta = -0.54, p < 0.001$ in alien identification) and its Interaction with Stimulus ($\beta = -0.48$ in suffix identification, $\beta = 1.00$ in alien identification, $p < 0.001$ in both cases). A stimulus effect was found only in suffix identification ($\beta = 0.62, p < 0.001$), not in alien identification ($\beta = 0.14, n.s.$).

Novel Stimuli We hypothesized that participants would apply the sociolinguistic association they had learned in the second training phase to novel words they had never seen before. The results show that identification with old and new words strongly mirrored each other in both conditions and both tasks. A mixed-effects model was fit on each of the two tasks, with Response (Suffix or Alien) as the dependent variable, Participant and Word as random factors, and Condition, Stimuli (either Alien or Suffix) and Novelty as fixed effects. The results showed no significant Novelty effect in suffix identification ($\beta = 0.23, n.s.$) and alien identification ($\beta = 0.10, n.s.$). These results indicate that the acquired sociolinguistic association could be generalized to new lexical items, and that first-encounter facilitation applies to both familiar and unfamiliar words.

Summary Our prediction concerning first-encounter facilitation was supported. That is, participants in the *NoExposure* condition were more likely to acquire the association between *-nup* and the Norl species than participants in the *Exposure* conditions, suggesting that the first encounter with a novel variant facilitated the acquisition of sociolinguistic variants of that variant. We also found that this effect extended to previously unseen words.

Experiment 2: Constraint violation

Experiment 2 used a similar paradigm to Experiment 1, but we modified the suffixation patterns to investigate a different source of surprisal. Instead of surprisal caused by encountering a variant for the first time, Experiment 2 investigated whether surprisal caused by encountering a linguistic variant in an apparently ungrammatical context (i.e., where it violated a grammatical constraint) would also facilitate the acquisition of sociolinguistic associations. We will term this constraint-violation facilitation.

Method

Participants 103 participants completed Experiment 2 online within 1.5 hours. After excluding participants whose duration was below the 2.5% quantile or above the 97.5% quantile of all participants, there were 97 participants left whose data were used for the final analysis. They were 69 females and 28 males, aged 17–78 (mean: 29.3) years. 28 of them were recruited from the University of Pennsylvania subject pool (and rewarded with course credit), and the remaining 69 were recruited through the Prolific Academic website (and paid \$5 each). There were 48 participants

in the *Conditioned* condition and 49 in the *Unconditioned* condition.

Materials and Procedure The same words and aliens were used in Experiment 2 as in Experiment 1. The procedure was also the same, consisting of two learning phases and a test phase with two tasks.

Experimental Conditions and Predictions There were two between-subjects conditions based on the linguistic environment for the suffix *-nup*, which is shown in Fig. 8.

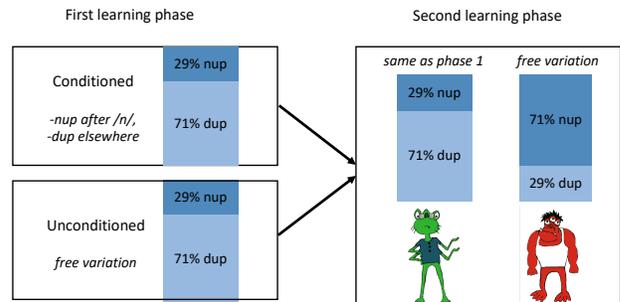


Figure 8: Variant distribution in the learning phases of Experiment 2

In the *Conditioned* condition participants in the first learning phase only ever saw *-nup* attached to the four nasal-ending stems (i.e., *falon*, *hiwen*, *tugan* and *wukin*), while *-dup* was attached to the 10 stems that did not end in a nasal. This implied a grammatical constraint on the distribution of *-nup* (i.e., that it only occurs after nasals). By contrast, participants in the *Unconditioned* condition were exposed to the two suffix variants in free variation (i.e., both *-nup* and *-dup* occurred with both nasal and non-nasal stems), though the variants still occurred at a ratio of ten (*-dup*) to four (*-nup*) – or 71% to 29% – just as in the *Conditioned* condition. In the second learning phase, Gulus exhibited precisely the suffixation pattern of the first phase, whereas Norls used the two suffixes freely across contexts at a ratio of four (*-dup*) to ten (*-nup*).

Similar to Experiment 1, we predicted that participants in the *Conditioned* condition would experience greater surprisal when they saw Norls using the two variants, especially *-nup*, in an ungrammatical way, and would be facilitated by this surprisal in learning the association between *-nup* and the Norl species, compared with those in the *Unconditioned* condition.

Results

On average, participants took 51 minutes ($sd = 12$) to complete the experiment and achieved a mean score of 152 ($sd = 13$).

Fig. 9 shows the aggregate results for suffix identification (left) and alien identification (right). The results do not appear to exhibit the predicted between-group difference in learning.

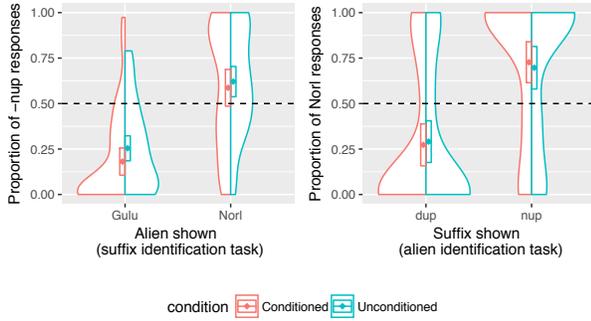


Figure 9: Proportion of *-nup* responses in suffix identification and Norl responses in alien identification (including 95% confidence interval). Dotted line indicates chance level.

A mixed-effects logistic model with Participant and Word as random factors, and Condition, Alien and their Interaction as independent variables revealed a significant Alien effect (Norl as the default, $\beta = -2.08, p < 0.001$) and a significant interaction ($\beta = 0.28, p = 0.012$), but no effect of Condition ($\beta = 0.28, n.s.$). In alien identification, a mixed-effects logistic model showed a significant effect of Suffix (*-nup* as the default, $\beta = 1.96, p < 0.001$), but no effect of Condition ($\beta = 0.06, n.s.$) and the Interaction ($\beta = -0.18, n.s.$).

Learning Proficiency It is possible that the absence of facilitation in Experiment 2 was due to variation in learning performance. The predicted facilitation depends on surprisal due to the apparent violation of a grammatical constraint. It therefore seems *a priori* clear that our predicted effect should occur only if participants learned the grammatical constraint. If they did not, violation of the constraint should not generate surprisal. To evaluate this possibility, we conducted a post hoc analysis in which we took participants' scores in the learning phase as a proxy for their learning performance. In particular, we divided participants into *good learners* and *poor learners* within each condition, according to whether their score was above or below the group mean. We then investigated whether constrain-violation facilitation could be found among good learners but not poor learners.

Fig. 10 shows the results for the 47 *good* and 50 *poor* learners. First, good learners showed a higher *-nup* rate for Norls and a lower *-nup* rate for Gulus in suffix identification, as well as a higher Norl rate for *-nup* and a lower one for *-dup*, compared with poor learners, indicating a better alignment between their responses and the pattern in the input, compared with poor learners. Second, the predicted learning facilitation is exhibited in the results of good learners, in that participants in the *Conditioned* condition exhibited a lower *-nup* rate for Gulus in suffix identification, and exhibited a higher Norl identification rate for *-nup* words and a lower Norl rate for *-dup* words in alien identification, compared with those in the *Unconditioned* condition.

We fit a mixed-effects logistic regression respectively on

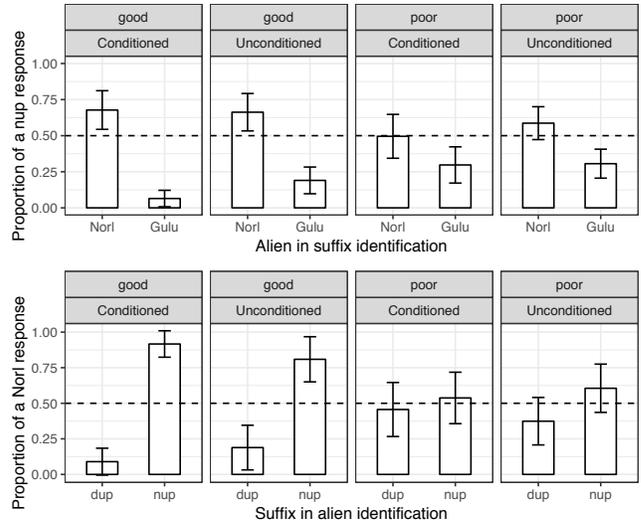


Figure 10: Proportion of *-nup* responses in suffix identification (top) and Norl responses in alien identification (bottom) by *good* and *poor* learners (including 95% confidence interval) Dotted line indicates chance level.

the data of good learners and poor learners. For good learners, the results showed significant effects of Alien ($\beta = 4.5, p < 0.001$), Condition ($\beta = 1.9, p < 0.001$) and their interaction ($\beta = -1.99, p < 0.001$) in suffix identification, as well as significant effects of Suffix ($\beta = 5.28, p < 0.001$), Condition ($\beta = 1.16, p < 0.001$) and Interaction ($\beta = -2.27, p < 0.001$) in alien identification. For poor learners, however, the results of suffix identification showed a main effect of Alien ($\beta = -0.93, p < 0.001$) and significant interaction between Alien and Condition ($\beta = -0.38, p = 0.008$), but no main effect of Condition ($\beta = 0.42, n.s.$). In alien identification, both factors of Suffix ($\beta = 0.33, p = 0.012$) and Condition ($\beta = -0.34, p = 0.007$) are significant, as is their interaction ($\beta = 0.62, p < 0.001$). Interestingly, however, the learning difference associated with the Suffix factor is the opposite of what was predicted: Learners in the *Unconditioned* condition did a better job in associating Norls to *-nup* than those in the *Conditioned* condition.

Novel Stimuli In evaluating whether acquisition effects were generalized to new words, we examined good and poor learners separately given their different patterns in acquisition. The results showed that although learners with different performance showed distinct patterns from each other, the behaviors with seen and unseen stimuli were highly consistent within each of the two learner groups. Good learners showed the correct alien-language association as well as additional facilitation from rule violation with both old and new words. Poor learners also showed consistent behaviors across old and new words, although behavior was mostly near chance level. Two mixed-effects models, one fit on each task, with Response as the dependent variable, Participant and Word as random factors, and Condition,

Stimuli, Novelty and their interactions as mixed effects, showed no significant Novelty effect or Novelty-relevant interactions.

Summary There was no evidence for constraint-violation facilitation in the aggregate results. However, post hoc analysis revealed that there was such an effect among “good learners”, participants who performed above the mean in training. This is consistent with the hypothesis, as constraint violation should facilitate learning only among individuals who have learned the constraint. Finally, the results of Experiment 2 replicate those of Experiment 1 in showing an ability to generalize acquired patterns, whether accurate or inaccurate, to new words.

General Discussion

We hypothesized that violation of expectation would cause a linguistic variant to be more salient and, as a result of this, that an association between this variant and a particular social group would be easier to learn. We tested this hypothesis in two experiments, each investigating a different kind of expectation violation.

The first experiment investigated exposure to a previously unencountered variant while the second investigated exposure to a variant that had previously occurred within a narrower grammatical context. In the first experiment the expectation violation had the predicted effect: Participants were more likely to associate the new suffix with the correct alien species (and the correct alien species with the new suffix) when the suffix had not been encountered in the initial learning phase. We also found that this effect extended to previously unseen words.

In the second experiment, we found the predicted effect, but only for *good learners*. While the division of Experiment 2 participants into good and poor learners was not planned and should therefore be taken with caution, the distinction has a clear precedent in earlier work (Rácz, Hay, & Pierrehumbert, 2017) and makes good theoretical sense. We should not expect violation of a grammatical rule to be salient to participants who have not learned that rule. Indeed, it would have been inconsistent with our hypothesis if we had found such an effect for participants who had not learned the rule.

Taken together, our results suggest that unexpectedness increases the salience of variants and makes their social distribution easier to learn, deepening our understanding of the role of individual language experience in the acquisition of sociolinguistic meaning.

References

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48.

Foulkes, P., & Docherty, G. (2006). The social life of phonetics and phonology. *Journal of phonetics*, 34(4), 409–438.

Fraundorf, S. H., & Jaeger, T. F. (2016). Readers generalize adaptation to newly-encountered dialectal structures to other unfamiliar structures. *Journal of memory and language*, 91, 28–58.

Hale, J. (2001). A probabilistic early parser as a psycholinguistic model. In *Proceedings of the second meeting of the North American chapter of the association for computational linguistics on language technologies* (pp. 1–8).

Jaeger, T. F., & Weatherholtz, K. (2016). What the heck is salience? How predictive language processing contributes to sociolinguistic perception. *Frontiers in Psychology*, 7.

Kaschak, M. P., & Glenberg, A. M. (2004). This construction needs learned. *Journal of Experimental Psychology: General*, 133(3), 450.

Labov, W. (1972). *Sociolinguistic patterns*. Philadelphia: University of Pennsylvania Press.

Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, 106(3), 1126–1177.

McDonald, S. A., & Shillcock, R. C. (2003). Low-level predictive inference in reading: The influence of transitional probabilities on eye movements. *Vision Research*, 43(16), 1735–1751.

McRae, K., Spivey-Knowlton, M. J., & Tanenhaus, M. K. (1998). Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language*, 38(3), 283–312.

R Core Team. (2014). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from <http://www.R-project.org/>

Rácz, P., Hay, J. B., & Pierrehumbert, J. B. (2017). Social salience discriminates learnability of contextual cues in an artificial language. *Frontiers in Psychology*, 8.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423.

Squires, L. (2014). Social differences in the processing of grammatical variation. *University of Pennsylvania Working Papers in Linguistics*, 20(2), 20.

Wickham, H. (2016). *ggplot2: Elegant graphics for data analysis*. New York: Springer-Verlag. Retrieved from <http://ggplot2.org>