THE LEARNING TRAJECTORY OF PHONOLOGICAL OPACITY*

Danica Reid[1], Larissa Melville[2], and Ashley Farris-Trimble[1]
[1]Simon Fraser University, [2]University of British Columbia

1. Introduction

A number of studies in recent years have attempted to experimentally test the biases that influence learning of phonological interactions like feeding, bleeding, counterfeeding, and counterbleeding (Bermúdez-Otero 2003; Brooks et al. 2013; Ettlinger 2008; Nazarov and Pater 2017; Rasin et al. 2017). Transparent interactions (feeding, bleeding) result in surface-true outputs; the surface form is consistent with each of the individual rules, which (in theory) is hypothesized to make the forms easier to learn (Kiparsky 1973). Maximal utilization describes a situation in which both of two interacting rules apply (feeding, counterbleeding), also hypothesized to facilitate learning (Kiparsky 1968). Artificial grammar and computational modeling studies have found support for both of these learning biases, albeit in different contexts (Jarosz 2016; Kim 2012; Prickett 2019).

Experimental tests of these biases have typically used palatalization and vowel deletion as their test case because they can be combined to create feeding, bleeding, counterfeeding, and counterbleeding interactions. For example, Prickett (2019) conducted an artificial grammar study comparing learning of these two processes in each of the four interaction types. Four groups of learners were presented with forms in which each individual process occurred and forms where the processes interacted in one of the four ways. He found evidence for both the transparency and maximal utilization biases. Participants in the feeding and counterbleeding conditions, both of which maximally utilized the rules and thus provided more evidence for palatalization, performed better on the palatalization trials than those in the counterfeeding and bleeding conditions. On the other hand, participants in the transparent bleeding and feeding conditions performed better on the interacting trials, suggesting that surface-true application enhances learning of that specific subset of trials.

The results of Prickett (2019) suggest that the transparency and maximal utilization biases affect different components of learning. The maximal utilization bias reflects the fact that the more evidence learners receive for a process, the better they learn it, whereas the transparency bias facilitates the acquisition of rule ordering itself. However, the processes used may have affected the learning process. The vowel deletion so commonly used in artificial language learning experiments that test transparent and opaque interactions (Prickett 2019; Jarosz 2016; Kim 2012) can result in the loss of creative information. This results in the presence of homophonous forms in a language, which can in turn detract from learning (Mazzocco 1997; Yin and White 2018). Specifically, in both

* Thanks to the members of SFU's Phonological Processing Lab for help with experiment design and data collection. Thanks also to attendees of the 2022 annual conference of the Canadian Linguistic Association for their suggestions and feedback.

Actes du congrès annuel de l'Association canadienne de linguistique 2022.
Proceedings of the 2022 annual conference of the Canadian Linguistic Association.

the bleeding and feeding interactions, the learners have to contend with homophony between the surface form resulting from the interacting processes and one of the other forms in the language (see (1) for the rules used and Table 1 for bleeding and feeding conditions from Prickett 2019). The current study uses a different rule interaction to test the transparency bias and examines not only the outcome but also the trajectory of learning.

(1) Rules from Prickett (2019)
Palatalization [t d] → [tʃ dʒ] / __ [i]
Deletion $V_1V_2 \rightarrow V_2$

Table 1. Trial types for bleeding and feeding languages in Prickett (2019). Homophonous forms are bolded.

		faithful	deleting	palatalizing	interacting
bleeding	UR	/t-a/	/k-i-a/	/t-i/	/t-i-a/
	SR	**[ta]**	[ka]	[tʃi]	**[ta]**
feeding	UR	/t-a/	/k-a-i/	/t-i/	/t-a-i/
	SR	[tʃa]	[kai]	**[tʃi]**	**[tʃi]**

1.1 Current Study

As part of a larger series of studies that look at how speakers learn and process multiple interacting rules, this study presents a partial replication of Prickett (2019) with particular focus on rule interactions that maximally utilize both rules. By looking at both processes, we hope to determine whether processing difficulties seen in previous experiments are due to the number rules applied to a word or to other factors (Farris-Trimble and Tessier 2019). In order to avoid the homophonous forms created by deletion rules, a vowel harmony rule was used instead of vowel deletion. The two rules are provided in (2) and (3).

(2) *Palatalization*
s → ʃ / __ {i/e}; z → ʒ / __ {i/e}
alveolar fricatives become post alveolar preceding front vowels

(3) *Vowel Harmony*
V → [α back, α round] / __ [α back, α round]
vowels assimilate to the backness and roundness features of a following vowel

Artificial language learning experiments targeting the learning of vowel harmony have found that speakers can and do learn harmony patterns for features of height, roundness, and backness (Finley 2017, 2021). The combination of backness and roundness harmony presented in (2) is a pattern found cross-linguistically in languages like Turkish, Tuvan, Tunica, and Ewe (see Rose and Walker 2011 for an overview of different attested harmony systems). By using palatalization and vowel harmony, each word form in the experiment is unique and non-homophonous. This allows for a determination of how the transparency and maximal utilization biases interact in languages like feeding and counterbleeding where both rules apply in interacting contexts. If the transparency bias holds across interactions that result from vowel harmony and not just deletion, then a transparent feeding language should be easier to learn than an opaque counterbleeding language.

2. Methodology

2.1 Participants

A total of 80 participants were recruited from the Simon Fraser University Linguistics department. All participants were students who received course credit for participation. Within our 80 participants, 48 were native speakers of English, with Mandarin, Cantonese, and Korean representing the largest proportion of non-English first languages. Once a participant signed up for the study, they were sent a link to the experiment which randomly assigned them to one of the two languages, counterbleeding or feeding, automatically. In total, 37 participants were assigned to the counterbleeding language and 43 were assigned to the feeding language.

2.2 Materials

A total of 72 base words were created for use across both the counterbleeding and feeding languages. Half of the words ended with /s/ or /z/ in order to undergo palatalization. The other half ended with /t/ or /d/ and did not undergo palatalization. Within each half of the base stimuli, half of the words were mono-syllabic and the other half were bi-syllabic. The initial consonants or initial syllables of the base forms were restricted to exclude any alveolar obstruents and, in each condition, only one word would start with any one consonant or initial syllable to maximize contrast.

Each language included three suffixes (diminutive, plural, and augmentative) that were attached to the base forms to present the environments for the palatalization and vowel harmony to occur. Both languages shared two of the same suffixes that appended in different orders to create the transparent or opaque interaction. The suffix /-i/ was used to represent the diminutive meaning and the suffix /-o/ was used to represent the plural meaning.

In the feeding language, diminutive plurals were created by appending the plural suffix /-o/ followed by the diminutive suffix /-i/. This triggered vowel harmony, resulting in the sequence [ei], which further triggered palatalization of a sibilant preceding the [e]. This resulted in a transparent surface form like [-ʃei]. The feeding language also contained a third suffix, /-e/, which represented the augmentative meaning and served to illustrate to participants that the vowel [e] resulted in palatalization, even when not derived by the vowel harmony rule. Examples of rule applications for the feeding language are shown in Table 2.

Table 2. Example of feeding rule interactions.

	/kas + o + i/	/kat + o + i/
Vowel Harmony	kas**ei**	kat**ei**
Palatalization	ka**ʃe**i	--
	[kaʃei]	[katei]

In the counterbleeding language, the diminutive plural was created by appending the diminutive /-i/ before the plural /-o/. Palatalization applied first, palatalizing a sibilant before the /-i/, which was then realized as [-u] because of the vowel harmony rule, resulting in an opaque output in which palatalization appears to have applied in the wrong context: [-ʃuo]. The counterbleeding language also contained a third augmentative suffix, /-u/, which was meant to ensure that participants understood that /-u/ did not trigger palatalization. Examples of rule applications for the counterbleeding language are shown in Table 3.

Table 3. Example of counterbleeding rule interactions.

	/kas + i + o/	/kat + i + o/
Palatalization	ka**ʃi**o	--
Vowel Harmony	kaʃ**uo**	kat**uo**
	[kaʃuo]	[katuo]

Each modified base form was presented in five different forms: the base, the diminutive, the plural, the augmentative, and the diminutive plural. A correct and foil form was recorded for each language. Foil forms included the suffixes applied in the reverse order. Auditory stimuli were recorded by the second author, a native speaker of Canadian English who could read IPA and knew the details of the experiment. Stimuli were recorded in Praat (Boersma and Weenink 2022) using a Maono A04 USB Microphone at 192kHz. Each word was recorded three times in the phrase "he said___" and a team of research assistants in SFU's Phonological Processing Lab went through to select the best token of each. In total there were 964 individual tokens selected for use across both languages.

In addition to the auditory stimuli, 72 images were selected for use in the experiment. Each of these images was edited to be presented in 5 possible forms: base, diminutive, plural, augmentative, and diminutive plural. The word-image pairs were randomly determined for each participant to control for any possible effect of word associations from other languages. A sample set of images is provided in Figure 1 to Figure 5.

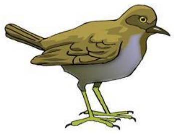

Figure 1. Base form

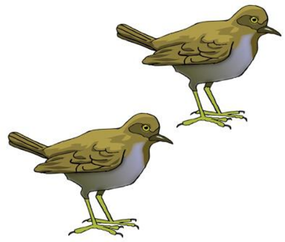

Figure 2. Plural form

Figure 3. *Diminutive form*

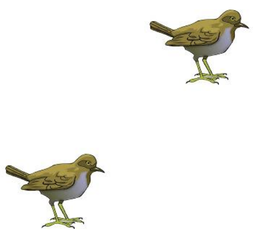

Figure 4. Diminutive plural form

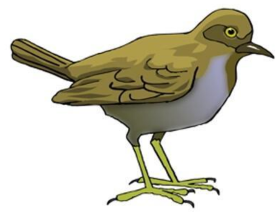

Figure 5. *Augmentative form*

2.3 Procedure

The experiment was fully coded in jsPsych (de Leeuw 2015) and distributed to participants through JATOS (Lange et al. 2015). It was divided into three parts: two training phases and a testing phase. In the training trials, participants would see the base image and hear its associated base word. They would then be shown a screen where the base image was presented with reduced visual saliency (it was presented as partially transparent) and the modified visual was presented with a red frame around it to indicate that this was the new target image. In the first training phase, participants were presented with four practice words that were not used in the remainder of the study. In order to help participants learn the suffixes associated with each modified form, this portion of training presented only the correct pronunciation of the modified word. Each base was presented four times so participants would hear the suffixes associated with each modification (plural, diminutive, diminutive plural, and augmentative) for a total of 16 trials. In the second phase of training, instead of only hearing the correct modified form, participants heard both a correct and foil version. For example, if the base form was /kas/ and the modified image showed the diminutive plural form in the feeding language, the two forms that they would hear for the modified auditory stimuli were the correct form /kaʃei/ and the foil form /kasei/. Throughout this training phase, participants were told to press the A key if they believed the first word was correct or the L key if they thought the second word was correct. It was made clear in the instructions that participants would probably be guessing in the early stages, but that they would be given feedback that would help them learn the patterns. After going through 108 training trials, participants were then moved into the testing phase where they were presented with words that they had not practiced in the training. This allowed us to determine how participants generalized their word learning. Training and testing trials followed the same procedure. A visual representation of the training and testing trials is provided in Figure 6. In total, participants went through 90 testing trials.

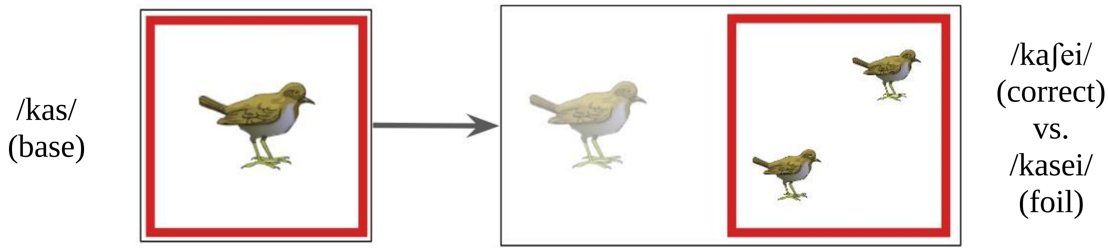

Figure 6. Visual representation of trial design.

Based on the predictions of the Transparency and Maximal Utilization biases and Prickett's (2019) findings, we expected both groups to perform well with the individual processes (palatalization and vowel harmony), and that differences between the feeding and counterbleeding groups would emerge primarily in the interacting trials. Here we expected the feeding group to outperform the counterbleeding group.

3. Results

We examined training trials visually to get a sense of the learning trajectory of each process, and then we used the participants' accuracy in the testing phase to evaluate our hypotheses.

3.1 Overall accuracy

We first analyzed participants' response accuracy, as the task we were asking them to do was not an easy one. Figure 7 shows test accuracy by participant across all trial types (90 test trials). Average accuracy across all participants in both language groups was 60%, *SD* = 16.1, with no difference between the feeding (*M* = 58.2, *SD* = 15.6) and counterbleeding (*M* = 62.1, *SD* = 16.7) groups, $t(75) = 1.06$, p = .29. The median accuracy was 55.6%, just above chance.

Because average accuracy varied so widely across participants, we looked closer at individual accuracy. Here we found that 53 of the 80 subjects were at or below 60% accuracy, with over a quarter of the participants (23) scoring lower than chance (50%). We refer to the 27 participants who scored greater than 60% as the high-accuracy group. The high-accuracy group consisted of 17 L1 English speakers (*M* = 78.6, *SD* = 12.4) and 10 L2

English speakers (*M* = 78.4, *SD* = 13.3), and there was no difference in accuracy between the two groups, $t(25) = 0.04$, p = .97).[1]

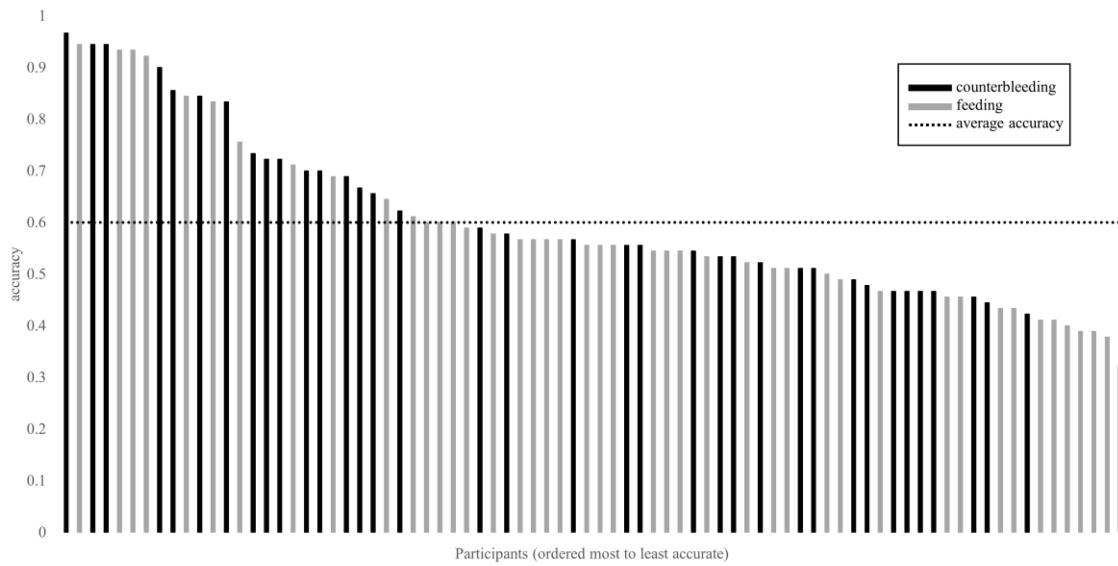

Figure 7. Individual participant accuracy; horizontal line shows average accuracy.

Rather than comparing learning of different processes in participants who did not learn much at all, we decided to limit our analysis to the 27 participants in the high accuracy group. All subsequent analyses include only that subset of participants. Because L1 and L2 English speakers performed similarly, we did not exclude participants on the basis of native language. In the discussion, we return to the question of why some participants had such difficulty with the task.

3.2 Training

Figure 8 shows the high-accuracy group's rolling accuracy average over a four-trial window. Note that the feeding group had more exposure to palatalization (because palatalization was fed by vowel harmony), while the counterbleeding group had more exposure to no-palatalization (because fewer words were in a palatalization context).

Following Prickett (2019), we expected both groups to learn vowel harmony well, because it was a process that applied transparently and without apparent exception. In the training data, we see that the two groups were accurate in vowel harmony trials from early on, with the feeding group perhaps improving some over time while the counterbleeding group remained stable.

---

[1] Indeed, there was no difference among L1 (*M* = 61.9%, *SD* = 16.6) and L2 English speakers (*M* = 57.3%, *SD* = 15.3) among participants at all accuracy levels, $t(78) = 1.24$, p = .217.

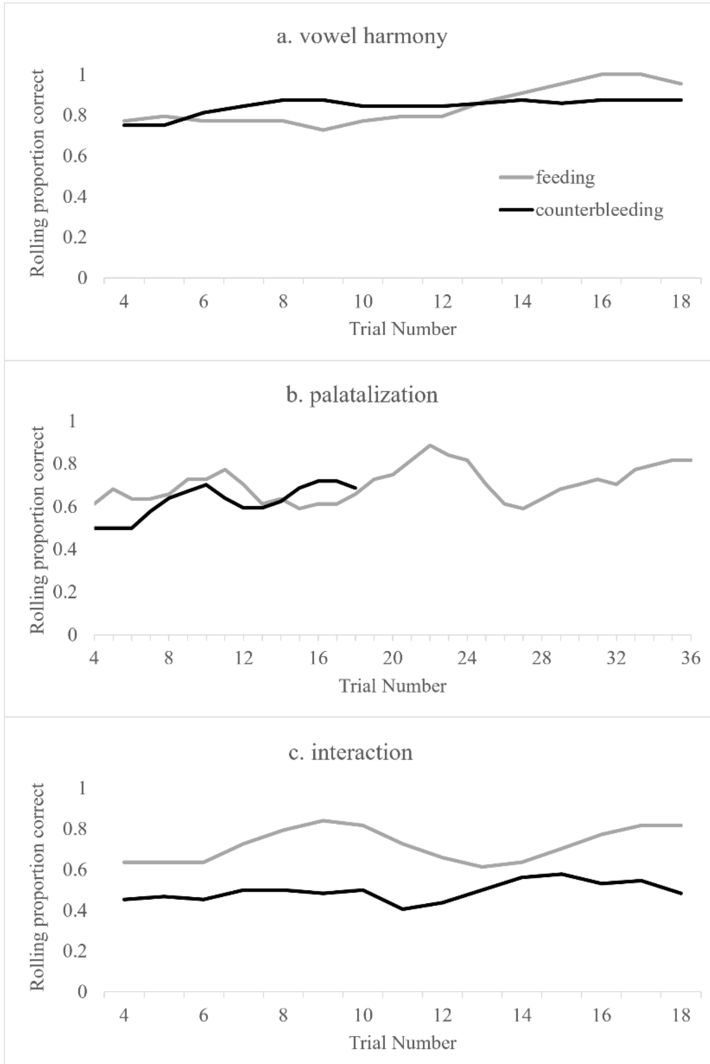

Figure 8. Rolling accuracy for high accuracy group in vowel harmony trials (a), palatalization trials (b), and interacting trials (c).

Prickett's (2019) findings suggest that a bias to maximize process utilization would mean that feeding and counterbleeding interactions, in which both rules apply, would perform equally well on palatalization trials (and outperform counterfeeding and bleeding patterns, which were not included in our experiment). Though training performance in the palatalization trials is lower overall than that of the vowel harmony trials, it does seem that the two groups performed similarly in the first 18 training trials. However, the feeding group was able to continue learning and reinforcing the palatalization pattern, while the counterbleeding group simply saw that pattern less frequently. By the end of training, the feeding group had increased in palatalization accuracy by about ten points relative to the counterbleeding group.

The clearest difference between the two groups comes from the interacting trials. In these trials, the feeding group outperforms the counterbleeding group from the start. The counterbleeding group hovers around chance performance, suggesting that they never really learned the interaction, while the feeding group performs with much higher accuracy.

3.3 Testing

Results for the testing phase of the experiment are shown in Figure 9. Both the feeding and counterbleeding groups achieved high accuracy in the vowel harmony trials. However, the feeding group was somewhat more accurate and considerably less variable in performance in both the palatalization and interacting trials. We ran a logistic mixed effects model to test these differences, with language group (feeding vs. counterbleeding) as a fixed factor and rule types dummy coded with two variables, palatalization and vowel harmony. Participants and items were included as random factors. We hypothesized that participants in the feeding group would out-perform those in the counterbleeding group, especially in interacting trials. In our model, a two-way interaction between language group and palatalization or vowel harmony would indicate that the groups performed differently on that type of trial; a three-way interaction between language group, palatalization, and vowel harmony would indicate that the groups performed differently in the interacting trials.

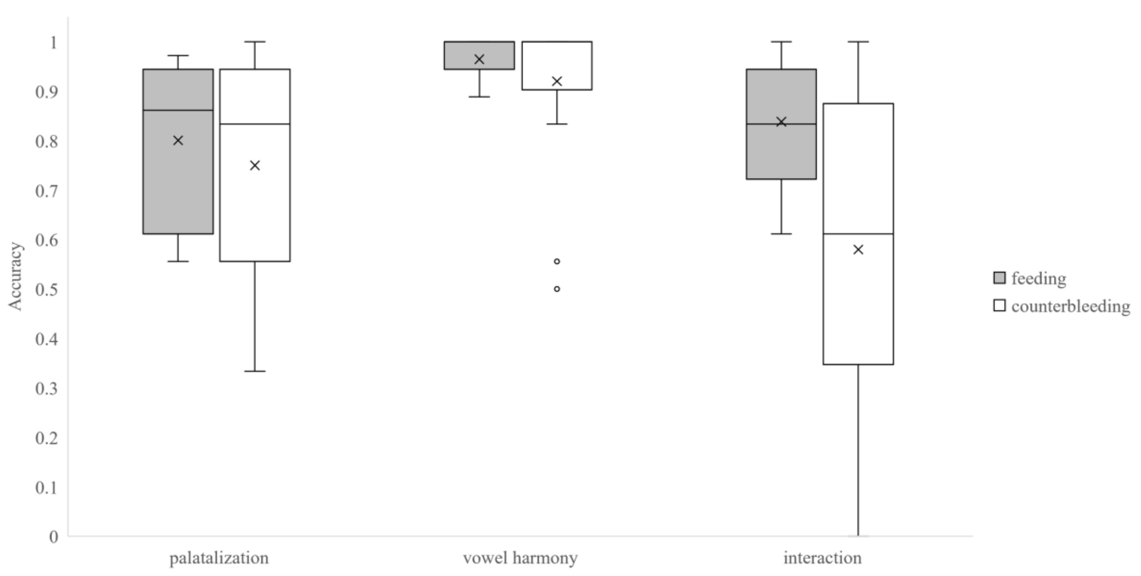

Figure 9. Test accuracy for the high accuracy group in the palatalization, vowel harmony, and interacting trials.

The interaction of language type with each individual rule reached significance (language group x palatalization: $z = 4.385$, $p < .001$; language group x vowel harmony: z

= 5.662, p < .001). The three-way interaction was marginally significant (language group x palatalization x harmony: $z = 1.671$, p = .0946). Collectively, these results indicate that the feeding group had an advantage in learning all three trial types. The large variability in performance in the counterbleeding group, particularly in the interacting trials, suggests that there may be other factors that influence whether a participant can learn the interacting pattern in the number of trials available.

4. Discussion and conclusion

This study, a partial replication of Prickett (2019), had several key findings: overall participant performance was quite poor, suggesting that the pattern was generally difficult to learn; participants learning a transparent interaction had an advantage over those learning an opaque interaction in all trial types; and interacting patterns involving vowel harmony produced results similar to those that involve vowel deletion. We will address each of these findings in turn.

We were surprised by the fact that only about a third of the 80 participants scored above 60% accuracy, while almost another third scored below chance (50% or lower). Though we expected that language background (L1 English vs. L2 English) might influence accuracy, it did not. To try to understand these results, we looked in greater detail at individual participant performance. It was certainly *possible* to learn the two patterns and their interaction: eight participants (four feeding, four counterbleeding) scored with 90% or greater accuracy. Moreover, accuracy on one pattern was not particularly predictive of accuracy on another; participant 80 (L1 English, counterbleeding) was 100% accurate on the vowel harmony trials, 83% accurate on the palatalization trials, and 0% accurate on the interacting trials. This participant learned each of the individual patterns, but, despite receiving training to the contrary, consistently selected the bleeding output in the interacting trials. Participant 83 (L1 English, counterbleeding) was 100% accurate in vowel harmony trials, but only 39% and 28% accurate in palatalization and interacting trials, respectively. And participant 51 (L1 English, counterbleeding) was 67% accurate in vowel harmony trials, 0% accurate in palatalization trials, and 78% accurate in interacting trials. This means that the participant consistently chose the unpalatalized form in the palatalization trials, but the palatalized form in the interacting trials. While we cannot put much weight on individual participant data, it does seem clear that some participants built phonologies that were not consistent with the training data.

Unfortunately, the low overall accuracy limits our interpretation of the results of the study. The participants who did learn supported our hypotheses, but the fact that so many did not learn raises a number of questions. This may have been an effect of the task or the processes we chose. Our study incorporated both phonological and lexical learning, which may have increased the cognitive load substantially, as participants had to learn phonological processes and suffix-to-meaning correspondences simultaneously. Perhaps with additional training or a different set of processes, participants would have learned more easily. But participants' overall poor performance also raises questions about how people represent interacting forms and processes. It is possible that some learners may have focused on learning specific forms but may not have had time to make larger generalizations about the phonological patterns. The participants who had difficulty generalizing interacting processes to new forms may support analyses of opaque interactions that rely on lexicalized opacity (e.g., Mielke et al. 2003; Sanders 2003).

In the high accuracy group, though, the participants who learned a feeding interaction outperformed those who learned a counterbleeding interaction in all trial types. This difference was small in the vowel harmony trials, for which both language groups received the same amount of evidence in training and approached ceiling in testing. In palatalization trials, the feeding group received twice the evidence in training as the counterbleeding group. This occurred because the three suffixes used in the feeding interaction were /-*e*/, /-*i*/, and /-*o*/, and both /-*e*/ and /-*i*/ triggered palatalization. In the counterbleeding interaction, though, the three suffixes were /-*i*/, /-*o*/, and /-*u*/, and only /-*i*/ triggered palatalization. This difference was necessary to provide each group with the evidence needed to deduce the processes. That is, in the feeding language, the backness harmony rule resulted in /-*oi*/ being realized as [-*ei*], and so it was important to include evidence that [e] *did* trigger palatalization. In the counterbleeding language, /-*io*/ was realized as [-*uo*], and we provided evidence that [u] alone *did not* trigger palatalization. It is inherent to the two sets of interacting processes that participants will receive different amounts of evidence for each, and that in turn influenced accuracy.

In the interacting trials, the feeding group (*M* = 83.8%, *SD* = 13.7) far outperformed the counterbleeding group, which was highly variable (*M* = 58, *SD* = 30.1). This replicates Prickett's finding of a transparency bias in interacting trials. The participants in the counterbleeding group chose the transparent foil (a bleeding interaction) nearly as often as they chose the correct opaque output.

Finally, we chose to avoid a deletion process to reduce homophony. Prickett (2019) found that participants were most accurate in the deletion rule, and that the feeding group outperformed the counterbleeding group, and these results were each replicated in our study. This is an important step in establishing that it is the interaction, rather than the specific processes involved, that drives the transparency bias.

While artificial grammar learning provides a window into the learning process, the next step is to examine real-world acquisition of phonological interactions. Preliminary results from our lab suggest that both children and adults have difficulty applying interacting processes that occur in their native languages to new forms in at least some tasks (Bondi et al. 2022). Additional research into the acquisition of interacting processes and their lexical representations will be necessary to uncover the biases and barriers that influence that acquisition.

References

Bermúdez-Otero, Ricardo. 2003. The acquisition of phonological opacity. In *Variation within Optimality Theory: Proceedings of the Stockholm Workshop on 'Variation within Optimality Theory'*, ed. Jennifer Spenader, Anders Eriksson, and Östen Dahl, 25-36. Stockholm: Department of Linguistics, Stockholm University.

Boersma, Paul and David Weenink. 2022. Praat: doing phonetics by computer. Version 6.2.23.

Bondi, Kira, Danica Reid, Sahibnoor Dhami, Allyson Ugalde, Sydney Walton, Yimin Wang, Anne-Michelle Tessier, and Ashley Farris-Trimble. 2022. Alienology: Learning Opaque Phonology. Paper presented at the 18th Conference of Laboratory Phonology.

Brooks, K. Michael, Bozena Pajak, and Eric Baković. 2013. Learning biases for phonological interactions. Poster presented at the 2013 Annual Meeting on Phonology, University of Massachusetts Amherst.

Ettlinger, Marc. 2009. Phonological chain shifts during acquisition: Evidence for lexical optimization. In *Proceedings of NELS 38*, ed. Muhammad Abdurrahman, Anisa Schardl, and Marin Walkow, 259–269.

Farris-Trimble, Ashley and Anne-Michelle Tessier. 2019. The effect of allophonic processes on word recognition: Eye-tracking evidence from Canadian raising. *Language* 95(1): e136-e160.

Finley, Sara. 2017. Locality and harmony: Perspectives from artificial grammar learning. *Language and Linguistics Compass 11(1):* 1-16.

Finley, Sara. 2021. Coarticulation and learnability of transparent vowels in vowel harmony. *Proceedings of the Linguistic Society of America* 6(1): 92-106.

Jarosz, Gaja. 2016. Learning opaque and transparent interactions in Harmonic Serialism. In *Proceedings of the 2015 Annual Meeting on Phonology*, ed. Gunnar Ólafur Hansson, Ashley Farris-Trimble, Kevin McMullin, and Douglas Pulleyblank.

Kim, Yun Jung. 2012. Do learners prefer transparent rule ordering? An artificial language learning study. *Chicago Linguistic Society* 48: 375–386.

Kiparsky, Paul. 1968. Linguistic universals and linguistic change. In *Universals in Linguistic Theory*, ed. Emmon W. Bach and Robert Thomas Harms, 170–202. New York: Holt, Rinehart and Winston.

Kiparsky, Paul. 1973. Abstractness, opacity, and global rules. In *Three dimensions of linguistic theory*, ed. Osamu Fujimura, 57-86. Tokyo: TEC.

Lange, Kristian, Simone Kühn, and Elisa Filevich. 2015. "Just Another Tool for Online Studies" (JATOS): An easy solution for setup and management of web servers supporting online studies. *PloS One,* 10(6), e0130834.

de Leeuw, Joshua R. 2015. jsPsych: A JavaScript library for creating behavioral experiments in a web browser. *Behavior Research Methods, 47*(1), 1-12.

Mazzocco, Michèle M. M. 1997. Children's interpretations of homonyms: A developmental study. *Journal of Child Language*, 24: 441–467.

Mielke, Jeff, Mike Armstrong, and Elizabeth Hume. 2003. Looking through opacity. *Theoretical Linguistics* 29:123–139.

Nazarov, Alexei. and Joe Pater. (2017). Learning opacity in Stratal Maximum Entropy Grammar. *Phonology* 34: 299–324.

Prickett, Brandon. 2019. Learning biases in opaque interactions. *Phonology* 36: 627-653.

Rasin, Ezer, Iddo Berger, Nur Lan, and Roni Katzir. (2017). Rule-based learning of phonological optionality and opacity. *Proceedings of NELS 48*: 269-282.

Rose, Sharon and Rachel Walker. 2011. Harmony systems. In *The Handbook of Phonological Theory*, Second Edition, ed. John A. Goldsmith, Jason Riggle, Alan C.L. Yu, 240-290. Blackwell.

Sanders, Nathan. 2003. *Opacity and sound change in the Polish Lexicon*. PhD dissertation, University of California, Santa Cruz.

Yin, Sora Heng. and James White. (2018). Neutralization and homophony avoidance in phonological learning. *Cognition* 179: 89-101.