

A morphosyntactic inductive bias in artificial language learning

Itamar Kastner (Humboldt-Universität zu Berlin) & Tal Linzen (Johns Hopkins University)

Recent work has explored the biases a learner might bring to the task of acquisition. Natural language input might contain many generalizations, not all of which are necessarily learned (Hayes et al. 2009; Hayes and White 2013; Becker et al. 2011). When two generalizations are available, is there an inductive bias to prefer one over the other? This study pitted a syntactic generalization against a phonological one in order to examine which biases guide the learner.

Generalizations. In some languages the trigger of a phonological alternation might be characterized in both phonological and other terms. For example, Tsez noun classes have semantic as well as phonological cues. When cues contrast in nonce words, Tsez adults and children assign noun class based on the phonological cues rather than the semantic ones (Gagliardi and Lidz 2014). Noun class is presumably encoded formally by a diacritic on the root or the nominalizer *n*. Within an artificial language paradigm, Culbertson et al. (2017) found that generalizations in noun class assignment rely on the salience of the cue; for example, animacy is a salient semantic feature, so it will be preferred to less salient semantic or phonological cues. While these studies of noun classes are instructive, they do not target the morphosyntactic features active in the language. Nouns have semantic content, but semantic distinctions are not necessarily related to the rest of the morphosyntactic system: a class diacritic has no deterministic semantic interpretation, whereas a syntactic feature like [PL] maps the noun onto a plural representation. **The current study** tested what generalizations learners make when presented with data in which a pattern has two potential sources, one which is (morpho-)syntactic and one which is (morpho-)phonological.

Allomorphic alternations in a language might be described in syntactic or phonological terms (e.g. Embick 2010). For example, in Modern Hebrew (right), the stem vowel alternates between /a/ and /i/: it is /a/ when the subject is 1st or 2nd person (morphosyntactic), but also when the suffix is C-initial (morphophonological). In our artificial language, participants were trained on input which was compatible with two generalizations: one surface-based, “phonological”, and one contentful, “syntactic”. They were then tested on held-out cases to see what generalizations were made (Wilson 2006).

	<i>hexnis</i> ‘inserted’	
	SG	PL
1	hexnas-ti	hexnas-nu
2M	hexnas-ta	hexnas-tem
2F	hexnas-t	hexnas-tem
3M	hexnis	hexnis-u
3F	hexnis-a	hexnis-u

Methods. Participants were presented with a language in which words were of the form CVC, or CVC-V and CVC-CV when suffixed. There were ten different words, assigned to ten different pictures, and four suffixes. The study consisted of three stages, presented visually, as in Fig. 1.

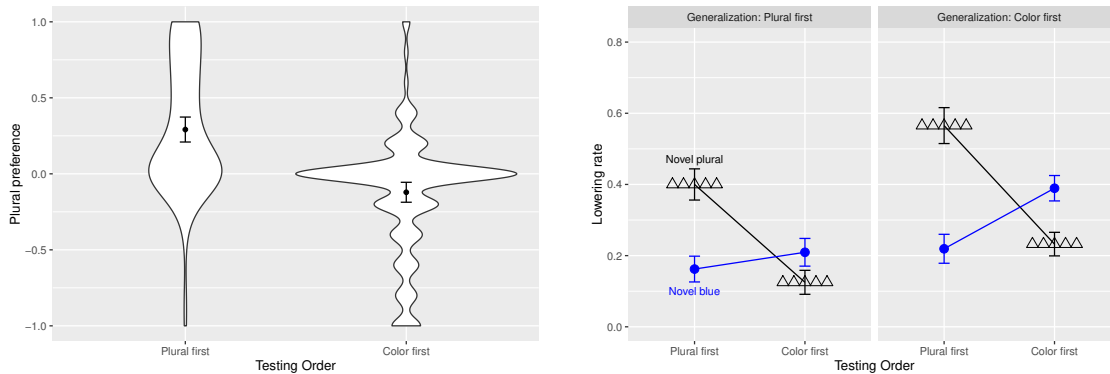
During Training, three words and two suffixes were introduced: V-initial for a red version of the shape and C-initial for a plural version. The plural suffix triggered lowering of the stem vowel to *-a-* (words in this stage had only non-low stem vowels). Participants were taught what each word means, (1)–(6) in Fig. 1. During Generalization, five new words were introduced. Participants were asked to pick the correct “red” and “plural” forms and feedback was provided after each trial, (7)–(10). This part ensured that the affixes were correctly generalized; participants who failed at least half the trials in this stage were excluded from analysis. Finally, during Test, two new suffixes were introduced: C-initial for a blue version of the shape and V-initial for a plural version, (11)–(15). These suffixes were demonstrated using two new words whose stem vowel was *-a-*, meaning they were ambiguous with regards to the cause of lowering. Participants were asked to pick the correct “blue” and “plural” forms from two choices, one faithful to the stem and another with lowering. The experiment was run online using Experigen (Becker and Levine 2013): 251 participants completed the study across four counterbalanced lists, and the data from 237 were retained after exclusions. Participants were recruited via Prolific Academic (prolific.ac).

Predictions. If participants made the “syntactic” generalization, they would lower the vowel before the “plural” V-initial suffix, based on the plural feature. If they made the “phonological” generalization, they would lower the vowel before the “blue” C-initial suffix, based on its phonological characteristics. Order of presentation within Generalization and within Test was counterbalanced across participants.

	New word introduced	Forced choice	Feedback
Training	(1) pil □		Given
	(2) pilu ■		Given
	(3) palti □□□□□□		Given
	(4) ter ○		Given
	(5) teru ●		Given
	(6) tarti ○○○○○○		Given
Generalization	(7) kun △		
	(8)	△△△△△△ kanti kunte	Given
	(9) bik ☆		
	(10)	★ biku biko	Given
Test	(11) mab ⇨		
	(12) mabgi ➡		
	(13)	★ bakgi bikgi	
	(14) mabi ⇨⇨⇨⇨⇨⇨⇨⇨		
	(15)	△△△△△△ kani kuni	

Figure 1: Experimental design.

Results. A main effect of Condition (lowering preference for novel plural over novel blue) was found in a mixed effects analysis ($p = 0.031$). A post-hoc analysis with Order of Testing as factor (plural first vs blue first) showed significant effects of Condition ($p = 0.007$), Order ($p = 0.035$) and their interaction ($p < 0.001$); see Fig 2a. Finally, an additional analysis adding Order of Generalization (plural first vs red first) revealed a main effect ($p < 0.004$) which did not interact with the other factors; see Fig 2b.



(a) Violin plots for subject preference, calculated as rate of plural lowerings minus rate of blue lowerings. (b) Interaction of lowering by condition with presentation orders. Error bars plot CIs within subjects.

Figure 2: Post-hoc analyses of presentation order during (a) Test, (b) Generalization and Test.

Discussion. Presented with ambiguous input that is equally compatible with a syntactic and a phonological generalization, do learners have a bias to prefer one to the other? Our results support the existence of a bias in favor of a morphosyntactic generalization, perhaps one sensitive to the existence of features such as [PL] in the grammar. In addition to lending initial support to a syntactic substantive bias, our results point to the crucial role of test item ordering in artificial language learning experiments. In future research, we plan to experiment with interleaved presentation of the two types of suffixes. We are also looking to extend this design beyond allomorphic interactions.