UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Can input statistics over-ride a prior bias in morpheme ordering? A test case with gender and number

Permalink

https://escholarship.org/uc/item/6n03n959

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 47(0)

Authors

Li, Chunan
Saldana, Carmen
Culbertson, Jennifer

Publication Date

2025

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Can input statistics over-ride a prior bias in morpheme ordering? A test case with gender and number

Chunan Li (chunanli@uchicago.edu)
Department of Linguistics
The University of Chicago

Carmen Saldana (carmencatalina.saldana@upf.edu)
Department of Translation and Language Sciences
Pompeu Fabra University

Jennifer Culbertson (jennifer.culbertson@ed.ac.uk)
School of Philosophy, Psychology, and Language Sciences
The University of Edinburgh

Abstract

In languages which mark both gender and number as distinct morphemes, there is a tendency to place gender closer to the noun stem than number. However, the typological data on this is sparse. Moreover, linguistic theories differ in how they explain ordering patterns of gender and number morphology: some theories focus on the structure of the representations of features in the speakers' minds, and other focus on the role of co-occurrence statistics. In a recent study, Saldana, Kanampiu, and Culbertson (in press) use artificial language learning to show that learners with a diverse range of language experience with grammatical gender and number exhibit a consistent bias for orders with gender closer to the noun stem than number. This order reflects the ordering in which most linguistic theories assume number and gender features are derived in word formation. Here, we build on this study to investigate how this bias interacts with the statistics of the linguistic input. In particular, we manipulate co-occurrence between stems and affixes so that learners are exposed to combinations of stems and number morphology more often than to stems and gender morphology. We test whether input statistics can push learners to reverse their natural preference, leading them to place number closer to the noun than gender. We find that our manipulation reduces, but does not reverse, the preference for gender-closest order. However, our study also highlights some difficulties learners have in acquiring novel features from sparse data. Ultimately, our findings highlight the dynamic interplay between representations of meaning and input-based learning mechanisms.

Keywords: gender; number; morpheme order; artificial language learning

Introduction

Morpheme order varies considerably across the world's languages, but certain patterns are more common than others. For instance, derivational morphemes, which alter a stem's category or meaning, typically attach closer to the stem than inflectional morphemes do (Bybee, 1985; Rice, 2011). Among inflectional morphemes, when distinct affixes for number and case appear together, number tends to be closer to the stem than case (Greenberg, 1963; Bybee, 2008).

The question of why these cross-linguistic trends emerge remains a topic of ongoing debate, with explanations ranging from language-specific cognitive constraints on representations, to domain-general mechanisms for chunking during learning and processing (Bybee, 2008; Chomsky, 1995; Evans & Levinson, 2009; Rooryck, Smith, Liptak, & Blakemore, 2010; Hay & Plag, 2004; Hahn, Degen, & Futrell, 2021). Recent studies have used artificial language learning experiments to explore the link between cognition and cross-linguistic ordering tendencies (e.g., Martin, Adger, Abels, Kanampiu, & Culbertson, 2024; Culbertson & Adger, 2014; Mansfield et al., 2022). For example, Saldana, Oseki, and Culbertson (2021) show that when participants learn two novel nominal affixes for number and case but are not told how they are ordered relative to one another, they tend to infer an order with number markers positioned closer to the noun stem than case markers. Along similar lines, Maldonado, Saldana, and Culbertson (2020) taught participants novel number and person verbal affixes, and found a preference for person morphemes to be placed closer to stems than number morphemes. In both cases, the authors argue that these findings reflect a general preference for compositionally transparent orders, in which representations of meaning, not input statistics drive order (Baker, 1985; Bybee, 1985; Rice, 2000; Culbertson, Schouwstra, & Kirby, 2020).

Here, we focus on the interaction between gender and number, which is more complex than e.g., number and case (Corbett, 1991, 2000), with sparse typological data on ordering. Previous experiments on gender and number ordering preferences have found that learners from a range of linguistic backgrounds, when taught novel gender and number affixes, but given no information about their order, reliably infer that gender should be closer to the noun than number (Saldana et al., in press). This suggests a robust cognitive bias favoring gender-closest ordering. The current study builds on this work to investigate how input statistics may interact with this cognitive bias. Below, we first briefly review relevant typological data and theories of the interaction between gender and number, and then turn to our study.

Typology and theories of gender and number order

Across the world's languages, gender systems are a common type of nominal classification device (Corbett, 2013; Kramer, 2015; Corbett, 1991). In many languages, gender is expressed through cumulative exponence with number, i.e., a single morpheme signals both gender and number features. However, in languages where gender and number are marked separately, gender typically appears closer to the noun stem (Fuchs, Polinsky, & Scontras, 2015). Despite the limited data on separative gender and number marking, some linguistic theories of gender and number are, in part, designed to predict these typological tendencies. For example, some theories argue for an abstract hierarchical representation in which gender and number are fused into a single morphosyntactic unit

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society* ©2025 the author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

(e.g., see Ritter, 1993; Carstens, 2003). These models do not make any predictions about relative order, however, they can help explain why the two features tend to be cumulative in languages. Alternative theories posit that gender and number occupy independent structural positions in noun derivation, with gender represented as part of the noun stem (e.g., see Alexiadou, 2004; Kramer, 2015; Panagiotidis, 2018), or projected outside the stem but below number in hierarchy (e.g., see Picallo, 1991; De Belder & Van Koppen, 2016; Antón-Méndez, Nicol, & Garrett, 2002). Assuming that linearisation reflects the order and locality relations in morphosyntactic composition (Rice, 2000; Baker, 1985; Bobaljik, 2012; Embick, 2010), these theories are supported by the tendency for gender to linearise closer to the noun stem than number.

While not specific to gender and number, other proposals argue that semantic relations between morphemes—i.e., how they modify or contribute to meaning—determine their linear order (e.g., Bybee, 1985; Rice, 2000). For example, morphemes with narrow semantic scope (either in general, or in a specific context) should be ordered closer to stems, while morphemes that take broader semantic scope should be ordered more peripherally (e.g., number morphology appears closer to the stem than case morphology; Bybee, 1985; Saldana et al., 2021). Applied to gender and number features, if focused on semantic scope alone (i.e., and not generally on morphosyntactic composition, cf. Baker, 1985; Bobaljik, 2012; Embick, 2010), these theories are only relevant to semantic gender, and in particular to animate nouns with variable gender. One could argue that number generally takes wider scope over gender, which itself takes narrow scope over the noun, as in *a herd$_{PL}$ of female$_{GEN}$ deer$_{N}$* (a *bevy* of *does*). The gender feature determines the specific meaning of the noun (i.e., a female rather than a male deer), and the number feature determines the numerosity of the set. If the scopal order were the reverse, the interpretation would be something like a "feminine plurality of deer". But this is unexpected, given that plurality is a formal property of sets which do not change their meaning based on gender features. Despite the tendency for morpheme order to reflect scope, some languages allow flexible ordering that does not affect meaning. An example comes from Chintang (Bickel et al., 2007), where three prefixes—*ma-* (negative), *u-* (third-person non-singular agent), and *kha-* (first-person non-singular object)—can appear in any of the six possible orders without changing the interpretation. This suggests that in certain languages, scope may not be the primary factor determining morpheme order.

In addition to explanations which focus on semantic or morphosyntactic representations, a third line of research attributes morpheme order to the distributional properties of linguistic input. This perspective argues that frequent co-occurrences between a stem and a particular morpheme can predict their relative closeness in linear order (Hay, 2001; Hay & Plag, 2004). Unlike approaches focusing on structural representations, this model emphasises usage frequency and the predictive relationships between linguistic elements as the primary determinants of morphological order (Hahn et al., 2021; Rathi, Hahn, & Futrell, 2022; Hahn, Mathew, & Degen, 2022). For example, Hahn et al. (2021) show that, in languages like Japanese and Sesotho, affixes that frequently co-occur with a verb stem tend to appear closer to it. This parallels findings at the phrase and sentence level, where the degree of dependency between heads and dependents has been argued to explain word order preferences (Futrell, Levy, & Gibson, 2020; Hahn, Jurafsky, & Futrell, 2020). Importantly, Hahn et al. (2021) clarify that this distributional approach does not necessarily oppose semantic or scope-based explanations. Instead, it may offer a way to operationalise these effects: if an affix has a stronger semantic impact on the stem, its distribution may be more restricted, leading to higher co-occurrence and closer placement. Therefore, these findings do not resolve whether semantic constraints are the underlying cause of morpheme order but indicate a potential link between semantic structure and distributional patterns (Saldana et al., 2021).

This kind of potential interaction is illustrated in the series of studies reported by Saldana et al. (2021), described briefly above. They argue that Greenberg's Universal 39, which states that number typically appears closer to the noun stem than case, reflects a cognitive preference for morpheme orders that align with semantic and morphosyntactic composition. Number impacts noun meaning more directly than case; the former modifies the noun's numerosity, while the latter connects the noun to broader event structure. In experiments where co-occurrence between stems and affixes was controlled—i.e., no distributional cues were available—participants inferred number-closest orders. However, in an experiment where the case allomorph was conditioned on the noun stem, the preference was attenuated and even reversed for some participants. These results illustrate that distributional statistics can sometimes override structure-based inferences.

Experimental evidence on gender-number order

To summarise, there are theories which hypothesise that gender and number are mentally represented such that gender is closer to the noun than number. These theories predict that learners should assume, all things equal, a linear order of these elements that reflects this relationship. This prediction appears to be borne out in the typology, however, in most languages with gender and number, these are expressed cumulatively. At the same time, there is evidence from cross-linguistic patterns, large-scale corpora, and experiments on case and number that suggest input statistics can also shape ordering preferences. To further investigate this, we build on the experimental paradigm used by Saldana et al. (in press) to test preferences for gender-number order in the absence of input cues. Saldana et al. (in press) trained adult speakers of English (a language with no grammatical gender), Italian (a language with suffixal cumulative gender-number inflectional morphemes), and Kîîtharaka (a Bantu language with pre-

fixal cumulative gender-number inflectional morphemes) on a novel gender-number paradigm, where nouns were assigned a gender value based on their semantics (e.g., female vs male entity, or animate vs inanimate entity). In the paradigm, singular number was unmarked, and plural was marked; similarly, one gender class was marked and the other unmarked. During training, instances which required the use of both gender and number markers were held-out. During the critical test, participants were shown the relevant meanings and had to guess how to combine the stem and two markers to convey the meaning in the language. Participants were taught either a language in which both markers were suffixes, or a language where both were prefixes. Across all three populations, and regardless of whether the markers were prefixes or suffixes, participants exhibited a strong preference for placing gender markers closer to the noun. This pattern aligns with formal theories which propose that gender is represented as distinct from number and occupies a hierarchically closer position to the nominal stem. It is also consistent with the idea that orders which reflect how a morpheme contributes to meaning composition are preferred.

Importantly for our purposes, in their experiments with Italian and English speakers in (Saldana et al., in press), co-occurrence statistics during training were balanced. In other words, each stem in the language was equally likely to occur with each affix marker. This indicates that a preference for gender-closest orders emerged even in the absence of biasing input. In this study we increase the salience of *stem + number* combinations to test the impact of co-occurrence statistics on ordering preferences. Our aim is to observe whether this shift could reverse or attenuate natural morpheme ordering preferences. If distributional cues can shift morpheme ordering preferences, we would expect to observe a preference towards number-closest orders by participants in our study. It is also possible that the preference for gender-closest orders is independent of, or strong enough to override conflicting cues from the input. If this is the case, we might expect to see results similar to those found in Saldana et al. (in press), pointing to a strong cognitive bias for gender-closest ordering.

Experiment 1

Design and materials

Following Saldana et al. (2021, in press), we trained participants on an artificial language whose nominal system had two distinct gender classes (animates vs. inanimates) and two numbers (singular vs. plural). Only one gender class and one number value were marked via affixation, allowing us to withhold instances where both gender and number affixes appeared simultaneously. This setup was intended to prompt participants during testing to infer the relative order of these morphemes, deciding whether to place the gender or the number morpheme closer to the noun stem. The two affixes were randomly chosen for each participant from the set {gu, sa, vi}.

In all conditions, the inanimate class was marked, while the animate class was not; plural number was marked, while singular number was not. As mentioned above, in Saldana et al. (in press), for English- and Italian speakers, in the input, all noun stems were equally likely to appear with both markers. This is achieved by using a single set of stems, which are shared across classes. For example the stem *shib* refers to a single (animate) sheep when it is unmarked, to a single (inanimate) ship when it is marked for inanimate, plural (inanimate) ships when it is marked for plural. Thus, the stem *shib* occurs equally often with no affix, with a gender class affix, and with a number affix. The critical test presents a plurality of (animate) sheep, where both markers are required.

In the current study, we use a similar design but increase the co-occurrence frequency of stem + *number* combinations. We achieve this by teaching them the plural forms for a different set of entities than those used to teach singular forms, and by having stems idiosyncratically shared across forms denoting inanimate and animate entities. In other words, particular stems always appear either in singular or plural form, and stems that appear in their singular form were homophonous across gender categories, and could denote animate (without a gender marker) or inanimate entities (together with a gender affix). This means that co-occurrence frequencies between each stem + number combination are more frequent than each stem + gender combination. There were 8 stems referring to four animate and four inanimate objects, adapted from Saldana et al. (in press). The full lexicon taught during the training phase is shown in Figure 1. Participants were randomly assigned to either a prefix or suffix condition. The affix order was held constant for each participant across all phases of training.

Singular		**Plural**
G1: animate	**G2: inanimate**	**G1: animate**
shib-Ø	shib-sa	lan-gu
weil-Ø	weil-sa	chit-gu
houf-Ø	houf-sa	dur-gu
kot-Ø	kot-sa	pek-gu

Figure 1: Training phase stimuli.

Importantly, in the original experiment (Saldana et al., in press), participants saw the same objects throughout training and testing. In other words, the held-out cases involved objects they had seen in training, just in a different numerosity. However, due to the design of our experiment, participants will necessarily see the same stems paired with new objects during the testing phase. The reason for this design change is that in order for some stems to occur only in plural form, these stems must show animate entities only, and cannot vary

in class (e.g., *lan-gu* meaning lions). The held-out test items are therefore inanimate plural entities corresponding to these stems (e.g., *lan-gu* meaning lamps). This change, which introduces new meanings in the critical test, may impact participants' behaviour. We will revisit this point in the discussion.

Experimental procedure

During the training phase, participants encountered three types of nominal forms: the stem alone (singular, unmarked class), gender-marked stems (singular, marked class), and number-marked stems (plural, unmarked class).

In the critical test trials, participants saw new meanings they hadn't seen before. These inanimate plural meanings are expressed by combining see stems with both affixes. See Figure 2.

Stem	Affix	
	Inanimate	Plural
lan	-sa	-gu
chit	-sa	-gu
dur	-sa	-gu
pek	-sa	-gu

Figure 2: Testing phase stimuli. Participants encounter unseen meanings which require seen stems plus both affixes.

The experiment was programmed using JsPsych (De Leeuw, 2015). Participants were instructed that they would be learning an "alien language" to describe pictures. Following Saldana et al. (in press), the training phase was segmented into three parts: Exposure, Picture-matching, and Recall.

During the Exposure phase, participants viewed images paired with nouns either alone or accompanied by a single affix indicating number or gender. See Figure 1: the stem alone (animacy.singular, where animacy $\rightarrow$ Ø, and singular $\rightarrow$ Ø), gender-marked stems (inanimate.singular, where inanimate $\rightarrow$ affix, and singular $\rightarrow$ Ø), and number-marked stems (animacy.plural, animacy $\rightarrow$ Ø, and plural $\rightarrow$ affix). Each stem in these three forms was shown five times, totaling 60 trials, randomly ordered for each participant.

In the Picture-matching phase, participants selected the correct image that matched the given nominal form. The incorrect image differed either in gender or number from the correct form. Participants encountered each pairing of correct and incorrect forms three times for 48 matching trials, randomly ordered for each participant. Feedback was provided after each trial.

The Recall phase tested participants on their learning of the forms encountered. They were presented with images and required to construct the corresponding form by clicking a set of buttons. These buttons displayed the correct stem, the gender affix, and the number affix, but were randomly ordered. Each stem, with different gender and number forms (animacy.singular, animacy.plural, and inanimate.singular), was shown three times across 36 trials, randomly ordered for each participant. Feedback was provided after each trial.

In the critical testing phase, participants were tested on the withheld form-meaning mappings for inanimate plural nouns. For these forms, both inanimate (gender) and plural (number) morphemes are overtly marked. Participants were required to actively construct the correct inflected forms by selecting a stem and two affixes from a set of buttons. The button set always showed the correct stem, the marked gender affix (inanimate), and the number affix (plural). These were randomly ordered on each trial. Participants had to use all three buttons and could not submit their answer until they did. Participants encountered each inflected stem three times, 12 trials in total, randomly ordered for each participant. No feedback was provided during these trials. See the procedure in Figure 3.

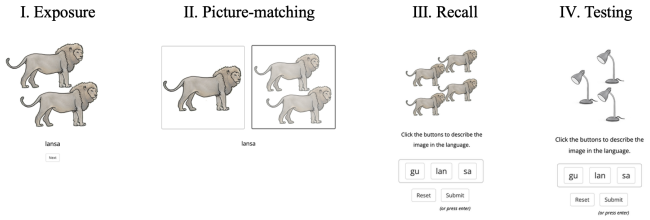

Figure 3: Full training and testing procedure.

Participants

52 monolingual native speakers of English were recruited from Prolific. The study took approximately 15 minutes to complete, and participants were compensated £2.25 each (an average payment rate of £9 per hour). Participants with accuracy below 75% in the Recall phase during training were excluded. After these exclusions (N = 1), the data of the remaining 51 participants were analyzed.

Data analysis

To analyse the participants' responses, we used Saldana et al. (in press)'s model structure and code available at https://osf.io/hmt7b/. We used R's *brms* package (Bürkner, 2018) as an interface to Stan (2017) to fit a mixed-effects Bayesian binomial model. This model predicts participants' production of the gender-closest order (e.g., "N-Gender-Number" if suffixing or "Number-Gender-N" if prefixing). Following Saldana et al. (in press), we also excluded trials that failed to place both affixes before or after the stem during testing.

Our dependent variable is participants' responses for each critical test trial, coded as 1 if the response was gender-closest ("N-Gender-Number" or "Number-Gender-N") and 0

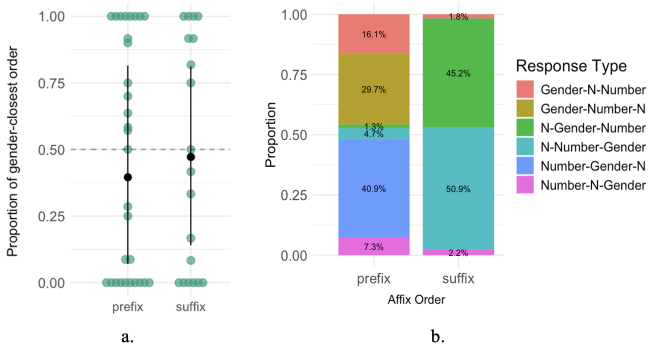

Figure 4: a. Proportions of gender-closest response order by different affix order in the current study. b. Proportions of different response types across affix order groups.

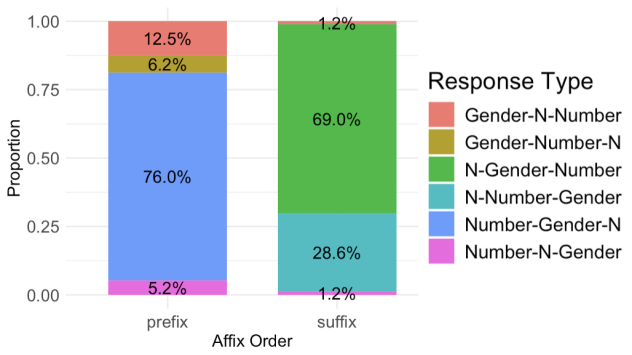

Figure 5: Proportions of response types for participants with correct number and gender understanding.

if it was number-closest ("N-Number-Gender" or "Gender-Number-N"). As fixed effects, we include affix order (suffix vs prefix condition, treatment-coded), and as random effects, we include intercepts for by-participants and meanings (i.e., objects). We used the priors set in Saldana et al. (in press): a Student-t distribution for the intercepts and fixed effects (DF = 6, $\mu = 0$, $\sigma = 1.5$), and a half-Cauchy prior with a scale parameter of 1 for the standard deviations of the random effects.

Results

A visual inspection of the results (see Figure 4a) suggest, unlike in Saldana et al. (in press), a shift away from any type of preference, be it for number-closest or gender-closest orders. The model's results confirm this observation. We find no evidence for a preference of number-closest orders (intercept: $\beta = -0.14$, 90% CI = [-2.14, 1.85], $P(\hat{\beta} < 0) = 0.55$), neither in prefixal nor in suffixal conditions (suffix vs prefix: $\beta = -0.41$, 90% CI = [-3.82, 2.93], $P(\hat{\beta} < 0) = 0.60$), indicating no evidence that suffixing versus prefixing affects the likelihood of producing gender-closest orders.

Figure 4b shows the proportions of various response types across different affix order groups, including responses where participants placed the stem between the two affixes. We also looked at the post-experimental survey. This indicated that 56.9% of participants reported correct interpretations for the number affix, 37.3% reported correct interpretations for the gender affix, and 29.4% reported correct interpretations for both. In the original experiments by (Saldana et al., in press), 70-75% of English and Italian participants (in animacy-based systems group) correctly interpreted the gender morphemes, and 85-100% correctly interpreted the number morphemes. Among Kîîtharaka-speaking participants, 93% correctly interpreted the number morphemes, and 38% correctly interpreted the gender morphemes. Compared to these previous experiments, the proportion of correct interpretations among our participants is relatively low.

Figure 5 shows the proportions of different response types among participants who correctly understand both number and gender morphemes. Unlike Figure 4b, these results show a preference for "gender-closest" answers. Given that only 29.4% of participants correctly identified the meanings of both affixes, we conducted a follow-up experiment to explore the potential preference for gender-closest responses among such participants.

Experiment 2

In Experiment 1, we observed more varied responses in the prefix order and more consistent responses in the suffix conditions, as in Figure 4b. For English speakers, suffixes align more closely with their native language structures. In Experiment 2, we replicated the procedures of the previous experiment but focused exclusively on the suffix order, examining the proportions of different responses among participants who correctly understood both number and gender morphemes. The materials and methods used in Experiment 2 were identical to those in Experiment 1, with the exception that only the suffix order was tested.

Results

We analysed the data using a similar model structure to that in Experiment 1, but without the fixed effects (i.e., an intercept-only model) as there was a single condition here. Consistent with the findings from Experiment 1, there is no evidence for neither gender-closest nor number-closest orders ($\beta = -0.57$, 90% CI = [-1.42, 0.27]), as shown in Figure 6a.

In Experiment 2, 32% of participants correctly identified both meanings. Figure 6b shows the proportions of different answers among these participants. The proportion of gender-closest answers is slightly higher, but there is no prominent preference compared to number-closest orders. The results indicate that the finding from Experiment 1—that participants who correctly understood both gender and number interpretations preferred gender-closest orders—was not replicated in Experiment 2.

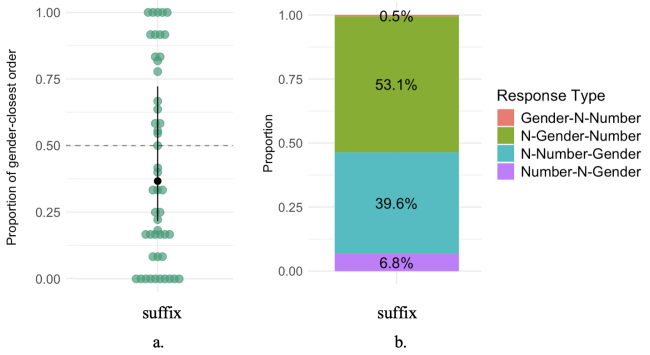

Figure 6: a. Proportions of gender-closest response order; b. Proportions of response types for participants with correct number and gender understanding

Discussion

This study aimed to investigate biases in morpheme ordering, specifically ordering of gender and number. Some theories of morpheme order predict that gender should be closer to the noun stem than number, however these theories make different predictions about the role of input statistics. Here we built on previous results which used artificial language learning experiments to test whether learners have a robust prior preference for gender-closest orders (Saldana et al., in press) in the absence of distributional cues from the input. We used this paradigm to explore whether participants' preferences for placing gender morphemes closer to the noun stem could be shifted by such cues. Specifically, we tested whether the preference was attenuated when number markers co-occur more frequently with particular noun stems than gender markers.

In Saldana et al. (in press), the frequency of co-occurrence for all stem + affix combinations was balanced for English and Italian speaking participant groups. Under these conditions, a robust preference for gender-closest orders was found across for both prefix and suffix systems and for both sex-based and animacy-based gender classification systems. In our study, we used a similar type of design, but skewed the distribution of stem-affix co-occurrences: one set of stems always occurred with a plural number marker, while another set of stems occurred half of the time with a gender marker and half without it. We found that this manipulation reduced participant's preference for gender-closest orders relative to Saldana et al. (in press), however, it did not reverse it—i.e., we did not observe a preference for number-closest orders, rather participants were at chance. This result suggests that the influence of input statistics on morpheme order preferences may not be as strong as biases based on structural representations. However, this interpretation of our findings was potentially undermined by the results of our post-experimental survey, which revealed that only a minority of participants correctly interpreted both gender and number affixes. We therefore separately analysed the data from participants who understood the affixes. Among these participants, there was a slightly higher proportion of gender-closest orders. We attempted to replicate this in Experiment 2. The results again revealed a reduced preference for gender-closest orders, with participants exhibiting at chance behaviour, this time regardless of whether they understood their meanings or not.

Our findings are thus consistent with the hypothesis that co-occurrence statistics can influence morpheme order (although this preference fails to over-ride a prior representational bias): after increasing the prominence of stem + number combinations, the gender-closest bias appears attenuated. However, it is also possible that the specific design of our task may have led to participants' at-chance behaviour. Specifically, the new experimental design may have been more challenging for participants compared to Saldana et al. (in press). In the critical test phase, our participants had to produce morphemes for stems they had seen, but with inanimate meanings they had *not* previously encountered. Thus, in addition to any influence that the manipulation of stem + number combinations may have had, the participants' lack of familiarity with these meanings may simply have increased their uncertainty. This interpretation is supported by the fact that participants showed lower understanding of morpheme meanings and also produced a wider variety of response orders than in Saldana et al. (in press). This points to the need for further research to differentiate between the effects of experimental design and the influence of stem + number combinations in future research.

Conclusions

This study investigated biases in morpheme ordering within nominal systems, focusing on English speakers' inferences about how to order gender and number morphemes in a new language. While there is some typological evidence that gender tends to appear closer to the noun stem than number, and this is built into at least some theories of morphosyntax, it is unclear what ultimately shapes morpheme ordering preference in this domain. Recent experimental work has aimed to test whether there is indeed a cognitive preference for this order. Saldana et al. (in press) found robust evidence for a preference for gender-closest orders in the absence of distribution cues that might have pushed learners toward one order or another. However, previous work looking at case and number morphology suggested that distributional cues may also play a role (Saldana et al., 2021). Through a series of experiments modeled after Saldana et al. (in press), we examined whether the effect of stem + affix co-occurrence statistics would be strong enough to reverse the apparent cognitive bias for gender-closest order.

While our manipulation did lead to a diminished preference for gender-closest orders among participants, the experimental design may also have posed particular challenges. Thus, while the possibility remains open that co-occurrence frequencies between stems and morphemes might affect learners' inferences about morpheme order, further work is needed to rule out alternative explanations in this case.

References

Alexiadou, A. (2004). Inflection class, gender and DP internal structure. *Explorations in nominal inflection*, 21–50. doi: 10.1515/9783110197501.21

Antón-Méndez, I., Nicol, J. L., & Garrett, M. F. (2002). The relation between gender and number agreement processing. *Syntax*, *5*(1), 1–25. doi: 10.1111/1467-9612.00045

Baker, M. (1985). The mirror principle and morphosyntactic explanation. *Linguistic Inquiry*, *16*.

Bickel, B., Banjade, G., Gaenszle, M., Lieven, E., Paudyal, N. P., Rai, I. P., ... Stoll, S. (2007). Free prefix ordering in chintang. *Language*, 43–73. doi: 10.1353/lan.2007.0002

Bobaljik, J. D. (2012). *Universals in comparative morphology: Suppletion, superlatives, and the structure of words* (Vol. 50). MIT Press.

Bybee, J. (1985). *Morphology. a study of the relation between meaning and form*. John Benjamins Publishing.

Bybee, J. (2008). Formal universals as emergent phenomena: The origins of structure preservation. In J. Good (Ed.), *Linguistic universals and language change* (pp. 108–121). Oxford University Press Oxford.

Bürkner, P.-C. (2018). Advanced Bayesian Multilevel Modeling with the R Package brms. *The R Journal*, *10*(1), 395–411. doi: 10.32614/RJ-2018-017

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of statistical software*, *76*. doi: 10.18637/jss.v076.i01

Carstens, V. (2003). Rethinking complementizer agreement: Agree with a case-checked goal. *Linguistic inquiry*, *34*(3), 393–412. doi: https://doi.org/10.1162/002438903322247533

Chomsky, N. (1995). *The minimalist program*. MIT Press.

Corbett, G. G. (1991). *Gender*. Cambridge University Press.

Corbett, G. G. (2000). *Number*. Cambridge University Press.

Corbett, G. G. (2013). Sex-based and non-sex-based gender systems (v2020.4) [Data set]. In M. S. Dryer & M. Haspelmath (Eds.), *The world atlas of language structures online*. Zenodo. doi: 10.5281/zenodo.13950591

Culbertson, J., & Adger, D. (2014). Language learners privilege structured meaning over surface frequency. *Proceedings of the National Academy of Sciences*, *111*(16), 5842–5847. doi: 10.1073/pnas.1320525111

Culbertson, J., Schouwstra, M., & Kirby, S. (2020). From the world to word order: Deriving biases in noun phrase order from statistical properties of the world. *Language*, *96*(3), 696–717. doi: 10.1353/lan.2020.0045

De Belder, M., & Van Koppen, M. (2016). One module, different levels of merge: An (n) compounds in dutch. *Studia Linguistica*, *70*(1), 1–33. doi: 10.1111/stul.12040

De Leeuw, J. R. (2015). jspsych: A javascript library for creating behavioral experiments in a web browser. *Behavior research methods*, *47*, 1–12. doi: 10.3758/s13428-014-0458-y

Embick, D. (2010). *Localism versus globalism in morphology and phonology*. MIT Press.

Evans, N., & Levinson, S. C. (2009). The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences*, *32*(5), 429–448. doi: 10.1017/S0140525X0999094X

Fuchs, Z., Polinsky, M., & Scontras, G. (2015). The differential representation of number and gender in spanish. *The linguistic review*, *32*(4), 703–737. doi: 10.1515/tlr-2015-0008

Futrell, R., Levy, R., & Gibson, E. (2020). Dependency locality as an explanatory principle for word order. *Language*, *96*(2), 371–412.

Greenberg, J. H. (1963). Some universals of grammar with particular reference to the order of meaningful elements. In J. H. Greenberg (Ed.), *Universals of language* (pp. 73–113). Cambridge, MA: MIT Press.

Hahn, M., Degen, J., & Futrell, R. (2021). Modeling word and morpheme order in natural language as an efficient trade-off of memory and surprisal. *Psychological Review*, *128*(4), 726. doi: 10.1037/rev0000269

Hahn, M., Jurafsky, D., & Futrell, R. (2020). Universals of word order reflect optimization of grammars for efficient communication. *Proceedings of the National Academy of Sciences*, *117*(5), 2347–2353. doi: 10.1073/pnas.1910923117

Hahn, M., Mathew, R., & Degen, J. (2022). Morpheme ordering across languages reflects optimization for processing efficiency. *Open Mind*, *5*, 208–232. doi: 10.1162/opmi_a_00051

Hay, J. (2001). Lexical frequency in morphology: Is everything relative? *linguistics*.

Hay, J., & Plag, I. (2004). What constrains possible suffix combinations? on the interaction of grammatical and processing restrictions in derivational morphology. *Natural Language & Linguistic Theory*, *22*(3), 565–596. doi: 10.1023/B:NALA.0000027679.63308.89

Kramer, R. (2015). *The morphosyntax of gender* (No. 58). Oxford University Press, USA.

Maldonado, M., Saldana, C., & Culbertson, J. (2020). Learning biases in person-number linearization. In *Proceedings of the 50th annual meeting of the north east linguistic society* (p. 163-176).

Mansfield, J., Saldana, C., Hurst, P., Nordlinger, R., Stoll, S., Bickel, B., & Perfors, A. (2022). Category clustering and morphological learning. *Cognitive Science*, *46*(2), e13107. doi: 10.1111/cogs.13107

Martin, A., Adger, D., Abels, K., Kanampiu, P., & Culbertson, J. (2024). A universal cognitive bias in word order: Evidence from speakers whose language goes against it. *Psychological Science*, *35*(3), 304–311. doi: 10.1177/09567976231222836

Panagiotidis, P. (2018). (grammatical) gender troubles and the gender of pronouns. In *Gender and noun classification*. Oxford University Press.

Picallo, M. C. (1991). Nominals and nominalizations in catalan. *Probus*, *3*(3), 279–316. doi: 10.1515/prbs.1991.3.3.279

Rathi, N., Hahn, M., & Futrell, R. (2022). Explaining patterns of fusion in morphological paradigms using the memory–surprisal tradeoff. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 44).

Rice, K. (2000). *Morpheme order and semantic scope: Word formation in the athapaskan verb*. Cambridge University Press.

Rice, K. (2011). Principles of affix ordering: An overview. *Word Structure*, *4*(2), 169–200. doi: 10.3366/word.2011.0009

Ritter, E. (1993). Where's gender? *Linguistic inquiry*, *24*(4), 795–803.

Rooryck, J., Smith, N. V., Liptak, A., & Blakemore, D. (2010). Special issue on Evans & Levinson's "The myth of language universals". *Lingua*, *12*(120), 2651–2758. doi: 10.1016/j.lingua.2010.04.012

Saldana, C., Kanampiu, P., & Culbertson, J. (in press). Gender comes first: Experimental evidence for the representation of gender closer to the noun stem than number across linguistic populations. *Linguistic Inquiry*. Retrieved from https://doi.org/10.31234/osf.io/xgjz2_v4

Saldana, C., Oseki, Y., & Culbertson, J. (2021). Cross-linguistic patterns of morpheme order reflect cognitive biases: An experimental study of case and number morphology. *Journal of Memory and Language*, *118*, 104204. doi: 10.1016/j.jml.2020.104204