

Speakers' cognitive representations of gender and number morphology shape cross-linguistic tendencies in morpheme order

Carmen Saldana (carmen.saldanagascon@uzh.ch)

Department of Comparative Language Science, University of Zurich, Switzerland
Center for the Interdisciplinary Study of Language Evolution, University of Zurich, Switzerland

Jennifer Culbertson (jennifer.culbertson@ed.ac.uk)

Centre for Language Evolution, University of Edinburgh, UK

Abstract

Languages exhibit a tremendous amount of variation in how they organise and order morphemes within words; however, regularities are also found. For example, gender and number inflectional morphology tend to appear together within a single affix, and when they appear in two separate affixes, gender marking tends to be placed closer to the stem than number. Formal theories of gender and number have been designed (in part) to explain these tendencies. However, determining whether the abstract representations hypothesised by these theories indeed drive the patterns we find cross-linguistically is difficult, if not impossible, based on the natural language data alone. In this study we use an artificial language learning paradigm to test whether the inferences learners make about the order of gender and number affixes—in the absence of any explicit information in the input—accord with formal theories of how they are represented. We test two different populations, English and Italian speakers, with substantially different gender systems in their first language. Our results suggest a clear preference for placing gender closest to the noun across these populations, across different types of gender systems, and across prefixing and suffixing morphology. These results expand the range of behavioural evidence for the role of cognitive representations in determining morpheme order.

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

Introduction

Languages exhibit a tremendous amount of variation in how they organise and order different categories of morphemes within words. Despite this variation, regularities are also found. For example, derivational affixes (which change a stem's category or meaning) tend to appear closer to the stem than inflectional affixes (which do not). Within inflectional morphology, when distinct affixes exist for number (e.g., plural) and case (e.g., accusative) and they appear together, number is ordered closer to the stem than case (Greenberg, 1963). There is a growing tradition in linguistics of explaining these kinds of ordering generalisations as the output of cognitive representations or biases active during language learning (e.g., Bybee, 1985; Saldana, Oseki, & Culbertson, 2021; Mansfield et al., 2022; Maldonado, Saldana, & Culbertson, 2020) and processing (e.g., Hawkins & Cutler, 1988; Gibson et al., 2013; Hahn, Degen, & Futrell, 2021; Hay, 2001). Here we follow this tradition to explore the linearisation of gender and number morphology.

Gender (or noun class) inflectional morphology categorises nouns into two or more classes, by definition reflected in agreement patterns on other elements (e.g., determiners,

verbs, auxiliaries or adjectives). Although nouns are sometimes assigned to classes arbitrarily, there is always a semantic core to these classes; nouns in a given class will tend to show some overlap in their semantic features. Animacy and social gender (or perceived biological sex in nonhuman animals) are the most common semantic cores in gender systems across the world's languages (Corbett, 2013). We refer to these systems as *animacy-based* and *sex-based* gender systems (though note they can co-exist; Kramer, 2015; Corbett, 1991).

Animacy-based systems classify human versus non-human entities, and/or animate vs inanimate entities, for example. In Swahili (see 1), there are a number of classes (18 in total), and class and number are indicated cumulatively in a single prefix (on nouns and also on agreeing elements). The class I (1a) morpheme is used for animate entities in the singular, the class II (1b) morpheme is used for those same entities in the plural. By contrast, classes III (1c) and IV (1d) are used for the singular and plural of a set of inanimate entities.

- (1) *Swahili (Atlantic-Congo)*
Animacy-based gender system
 - a. m-tu
I-person
'person'
 - b. wa-tu
II-person
'people'
 - c. m-ti
III-tree
'tree'
 - d. mi-ti
IV-tree
'trees'

Sex-based systems categorise nouns into what are often called feminine (FEM) and masculine (MASC) gender classes. For example, in Tamil (Dravidian), all MASC nouns denote male humans or deities (Corbett, 2000, p 9) and FEM nouns denote females, the rest are classified as neuter. In Romance languages the correspondence is less strict. Entities denoting females will generally be FEM (e.g., 2a), and males generally MASC, but inanimate nouns are also categorised as MASC or FEM (e.g., 2a-i-ii). In Italian, the noun's final vowel is a robust phonological cue to class when it cannot be inferred by semantics (e.g., most FEM nouns end in *-a* in singular and *-e*

in plural); MASC in *-o* and *-i* respectively). As in other Romance languages, Italian also uses the same endings to mark social or perceived biological gender in some animate nouns: the same stem with *-a* or *-o* denotes the female or male entity respectively (see 2b). These nouns are often referred to as “common-gender” or “same root” nouns (Kramer, 2015). We will refer to this phenomenon as *variable* gender marking in opposition to *fixed* gender marking (e.g., 2a or 1).

(2) *Italian (Indo-European)*
Sex-based gender system

a. Fixed gender:

- i. la barca
the.FEM.SG boat.FEM.SG
- ii. le barche
the.FEM.PL boat.FEM.PL
- iii. la madre
the.FEM.SG mother.FEM.SG
- iv. le madri
the.FEM.PL mother.FEM.PL

b. Variable gender:

- i. la cerv-a
the.FEM.SG deer.FEM.SG
- ii. le cerv-e
the.FEM.PL deer-FEM.PL
- iii. il cerv-o
the.MASC.SG deer-MASC.SG
- iv. i cerv-i
the.MASC.PL deer-MASC.PL

The relationship between gender and number morphology is potentially complex, and linguistic theories differ in how the two are formalised. On one approach, gender, unlike number, is treated as an inherent property of the nominal in one way or another (e.g., either located on N and thus part of the lexical entry, or on the nominaliser; Harris, 1991; Alexiadou, 2004; Mel’čuk, 2013; Kramer, 2015). An alternative approach does not consider gender to be located on the nominal (Carstens, 2003; Carminati, 2005; Picallo, 1991; Ritter, 1993; Antón-Méndez, Nicol, & Garrett, 2002). Within this approach there are two main models. One argues that gender is always bundled together with number (e.g., Ritter, 1993; Carstens, 2003). The other argues that gender and number are split (e.g., Carminati, 2005; Picallo, 1991; Antón-Méndez et al., 2002). Within this, some argue that gender is bundled with number for variable gender marking, but not for fixed gender (Carminati, 2005; De Vincenzi, 1999).

Each of these models capture different cross-linguistic tendencies regarding gender and number morphology. If the two are both inflectional features that are bundled together, this predicts that marking gender and number *cumulatively* within a single morpheme, should be common. Indeed, a survey of grammatical affixes in the AUTOTYP typological database (Bickel et al., 2022) suggests that when gender is marked via affixation, it tends to appear within the same affix as number, both on nominals and on verbs (see fig. 1). Italian is a good example of such a case of gender-number cumulative affixation in nominals (see 2). At the same time, models that treat number and gender as split (be gender an inherent part

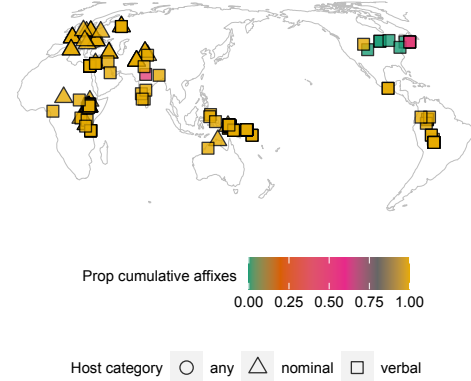


Figure 1: Proportion of cumulation in gender-number affixes in the AUTOTYP grammatical markers data (Bickel et al., 2022). The data shown includes 90 different languages from 37 different families.

of the nominal or not) predict another cross-linguistic tendency: In languages where number and gender morphology can be descriptively separated, gender morphology tends to be placed closer to the stem than number. An example of such language is, Spanish, which has a gender system similar to Italian, but with number morphology that can be described separately from gender. In Spanish (see 3), feminine nouns are typically indicated by a final vowel *-a* (3a), and masculine with *-o* (3b). Plural is then marked with the additional word-final affix *-s* (3c and 3d).

(3) *Spanish (Indo-European)*

- a. la barca
the.FEM.SG boat.FEM.SG
- b. el faro
the.MASC.SG lighthouse.MASC.SG
- c. las barcas
the.FEM.PL boat.FEM.PL
- d. los faros
the.MASC.PL lighthouse.MASC.PL

While the cross-linguistic data suggest that gender morphology tends to be placed closer to the stem than number, the evidence for relative order is scarce: there are not many languages with distinct (or separative) gender and number morphemes (see 1). It is therefore hard to assess whether alternative orders (i.e., stem-GENDER-NUMBER) are possible but simply missing from the small sample. If gender is part of the lexical entry and number is not across the board, gender should appear closer to the stem than number. However, if gender is always bundled with number (Ritter, 1993; Carstens, 2003) more variability in the linearisation of gender and number morphology is expected¹. Moreover, the pervasiveness of cumulative gender+number morphology and the frequent coexistence of fixed and variable gender within the same linguistic system, make it difficult to explore potential differences in the linguistic representations between them. If gender is derived with the noun stem in fixed gender but with number in variable gender (e.g., since in this case the stem

¹Though note that some of these theories posit an additional ad-hoc mechanism forcing number to dominate gender when they are bundled together (e.g., Ritter, 1993).

has no pre-specified gender (Carminati, 2005; De Vincenzi, 1999), we would expect more variability in the relative order of gender and number morphology in variable gender as compared to fixed gender.

Of course, even if natural language data were plentiful, cross-linguistic tendencies cannot be directly linked to individual-level representations or biases. These tendencies could be due to mere historical contingencies, or to distributional properties of the linguistic input. For instance, previous research has argued that co-occurrence statistics among stems and morphemes, rather than any abstract representations of their structure, may determine morpheme order and morphological fusion both within and across languages (Hahn et al., 2021; Hahn, Mathew, & Degen, 2022; Rathi, Hahn, & Futrell, 2022; Hay, 2001; Hay & Plag, 2004). Laboratory experiments allow us to test the impact of hypothesised representations on language learning directly. In recent work, (Saldana et al., 2021), used a series of artificial language learning studies to show that learners' (English and Japanese native speakers) inferences about number and case morphemes follow the cross-linguistic trend in having number closer to a noun stem than case—even when co-occurrence statistics among stems and different affixes were held constant. (Saldana et al., 2021) argue that this ordering preference therefore likely reflects a distinction between inherent inflection (here number) and contextual inflection (case) on the users' mind: Number marking is most relevant to the noun stem alone and thus more independent of the noun's role in the argument structure of a sentence, while case is fully dependent on it.

Here, we use a similar experimental design to (Saldana et al., 2021) to test whether the tendency to order gender closer to the noun stem than number reflects individuals' biases—in line with the representations hypothesised by the theories discussed above—active during language learning, and independent of co-occurrence statistics. We further explore whether this tendency is comparable across fixed and variable gender systems, or whether there are differences in the representations of fixed and variable gender marking which could lead to difference in the strength of that general bias. In particular, this could lead to alternative orders being more likely for variable gender systems than for fixed gender systems.

In our study participants are trained on miniature gender-number nominal paradigms. These paradigms instantiate either (animacy-based) fixed or (sex-based) variable gender. Participants' training input indicates whether affixes generally precede or follow the noun stem, but participants are not given any examples in which overt gender and number affixes co-occur on the same noun. At test, they are asked to produce an order for these held out examples. These productions will reveal the inferences that language learners make regarding the relative order of gender and number morphology, in the absence of any evidence in the input linguistic system. We run the same experiment across two different populations, English and Italian speakers, with substantially different gender systems in their first language (as we will further explain

in Participants). To preview, the order participants infer indicates a clear preference for placing gender closest to the noun, in accordance with a representation in which gender is more closely tied to the noun. We also find that Italian speakers, and not English speakers, are slightly more likely to produce the reverse orders for variable gender systems than for fixed gender systems. These results therefore expand the range of behavioural evidence for the role of cognitive representations in determining morpheme order, and further suggest a difference between the representations of gender across variable and fixed systems (consistent with the models proposed in , e.g., Carminati, 2005; De Vincenzi, 1999) worth exploring further in future work.

Materials and Methods

We ran two artificial language learning experiments with an extrapolation design to test participants' ordering preferences on gender and number affix ordering on nouns (following Saldana et al., 2021). In Experiment 1 we test native English speakers, and in Experiment 2, we test native Italian speakers. Participants were trained on a subset of a nominal system with two gender values (or classes) and two number values (singular and plural). In this system, only one gender value is overtly marked via affixation, the other is not; and only one number value is overtly marked via affixation, the other is not. The system is thus designed such that only some forms involve affixes of both gender *and* number. Crucially, these forms are held out during training, but participants are asked to infer them during testing. During training, participants therefore learn how single morphemes are ordered relative to the noun stem, and must therefore infer the relative order of the two morphemes at test. They can either place the gender morpheme closer to the noun stem than the number morpheme—as predicted—or vice versa.

Experiments were designed using a 2×2 between-subjects factorial design where we manipulate the type of affixation (suffixing or prefixing) and the type of gender system. In one condition, the type of gender systems marks the perceived biological sex of nonhuman animals with sexual dimorphism. In this type of system, the same stem is unmarked when the denoted animal shows female characteristics (FEM), and marked when it shows male characteristics (MASC). In this case, gender marking is uniquely encoded in the affix or zero marking, the stem in isolation cannot determine the gender of the noun. We refer to this condition as *sex-based*, but it is important to note that gender in this condition is crucially also variable because the same stem can be inflected for both genders. The other condition, by contrast, has a fixed gender system. In this condition, gender distinguishes between animate (AN) and inanimate (INAN) entities and the denoted entity type cannot be both AN and INAN. Stems denoting animates are marked with an affix, and stems denoting inanimates are zero-marked. However, unlike in the sex-based condition, the stems here do not share any semantic properties. We refer to this condition as *animacy-based*.

Artificial lexicon

Participants were trained on a language with a small lexicon of four stems, and two affixes (both prefixing or both suffixing depending on the condition), one expressing number (singular or plural) and the other expressing gender (sex-based or animacy-based depending on the condition). The two affixes were chosen randomly by participant from the set {gu, sa, vi} (Exp 1) or {gu, sa, di} (Exp 2). Examples of the full lexica for each gender system type is shown in Table 1. Gender and number affixes are attached to stems, with the relative order of affixes and nouns (prefixing or suffixing) determined by the condition. Plural number is marked via affixation, while singular is unmarked; similarly, one gender value is marked with an affix and the other is not. Feminine/inanimate gender is unmarked and masculine/animate is marked.

During training phases (described below), participants saw three kinds of nominal forms in the language: the stem alone (singular, unmarked class), gender marked stems (singular, marked class), and number marked stems (plural, unmarked class). Crucially, forms in which both a number and gender morpheme would be required—i.e., plural, marked class—were held out. Each of the four stems is used equally frequent in both classes, and with both numbers, therefore no co-occurrence statistics that differentiate stem+number from stem+gender combinations are present during training.

Experimental procedure

The experiment was programmed using *JsPsych* (De Leeuw, 2015) and displayed in participants' browser windows. At the start of the experiment, participants were told they would learn how to describe simple pictures in a new language. The training phase was divided into three sub-phases: exposure, picture-matching, and recall. During exposure, participants saw images along with orthographically presented nouns using a stem alone, or a stem with a single affix (either number of gender). More specifically, they were trained on three different types of inflected forms (see Table 1): the stem alone (G1.SG, where SG $\rightarrow \emptyset$ and G1 $\rightarrow \emptyset$), gender marked stems (G2.SG, where SG $\rightarrow \emptyset$ and G2 \rightarrow affix), and number marked stems (G1.PL, where PL \rightarrow affix and G1 $\rightarrow \emptyset$). Participants saw each stem in each of these three inflected forms five times, for a total of 60 trials (randomly ordered for each participant). In the picture-matching phase, participants saw two images and a form and had to click the corresponding image. As in the exposure phase, forms featured either a stem alone, or a stem with a single morpheme (either number of gender). The foil image for each trial was either the wrong gender, or the wrong number. Participants saw each stem with each pairing of correct and foil three times, for a total of 48 matching trials (randomly ordered for each participant). In the recall phase, participants were tested on the forms they had been trained on so far (i.e., G1.SG, G1.PL and G2.SG). In each trial, they saw an image and had to construct a corresponding form by clicking buttons. The button set always included the correct stem, the gender affix, and the number affix, but these

were randomly ordered on each trial. Full feedback was provided. Participants saw each stem inflected for G1.SG, G1.PL and G2.SG three times (36 trials in total, randomly ordered).

















In the critical testing phase, participants were tested on the held-out G2.PL form-meaning mappings: where both G2 gender and PL number are overtly marked, and therefore they need to provide forms with the two affixes. Participants saw an image and had to construct an inflected form by clicking buttons. The button set always included the correct stem, the gender affix (MASC or AN depending on the condition), and the number PL affix, but these were randomly ordered on each trial. Participants had to use all three buttons and could not submit their answer until they did. No other feedback was provided. Participants saw each inflected stem three times (total of 12 trials, randomly ordered).

Participants

















In Exp 1, we test English-speaking participants, whose native language has number marking in nouns, typically using a suffix (singular is unmarked, plural marked), but not gender marking. English only marks gender on third person singular personal pronouns, and does not have a noun class system. Moreover, pronouns in English are suppletive (i.e., inflectional values are non-segmentable), and thus participants have no evidence from their native language (neither for sex-based nor for animacy-based systems) of how inflectional morphemes of gender and number should be ordered relative to one another. They do, however, have evidence of derivational gender morphemes being placed closer to the root than (inflectional) number (e.g., *lion-ess-(e)s*), which in the absence of any other gender linguistic representation, could lead to participants placing sex-based gender markers closer to the stem because they treat them as derivational morphology. In Exp 2 we test Italian speakers in order to mitigate this potential confound. Italian, unlike English, has nominal inflectional morphology of both gender and number but it is marked cumulatively within the same affix, in other words, a single affix contains both the information of the gender and the number values of the noun (see 2 above). If Italian speakers also prefer to place number in the periphery, it will not be because that is the only inflectional affix they are familiar with, and we can conclude that it is more likely that this general preference is based on the relationships between morphemes and stems, independent of the participants' prior linguistic knowledge.

Participants were recruited through Prolific (www.prolific.co) and paid 2.00 GBP for participation in the study, which took approximately 12 minutes. We used Prolific screening criteria to include native English and Italian speaking participants who were raised monolingual and whose primary language is still their native language. As per our preregistration², participants who got less than 75% correct on recall trials in the training phase were excluded.

²All data and analyses reported are available at osf.io/5sjeb. The preregistered hypothesis, design and analysis plan can be found in osf.io/z2c9m (for Exp 1) and osf.io/un8yx (for Exp 2).

Exp 1 (English)			Exp 2 (Italian)		
SG		PL	SG		PL
G1:FEM	G2:MASC	G1:FEM	G1:FEM	G2:MASC	G1:FEM
 dur-0	 dur-sa	dur-gu	 cheru-0	 cheru-sa	cheru-gu
 lan-0	 lan-sa	lan-gu	 lonu-0	 lonu-sa	lonu-gu
 pek-0	 pek-sa	pek-gu	 povu-0	 povu-sa	povu-gu
 chit-0	 chit-sa	chit-gu	 kalu-0	 kalu-sa	kalu-gu

(a) Sex-based gender system.

Exp 1 (English)			Exp 2 (Italian)		
SG		PL	SG		PL
G1:INAN	G2:AN	G1:INAN	G1:INAN	G2:AN	G1:INAN
 shib-0	 shib-sa	shib-gu	 peta-0	 peta-sa	peta-gu
 kot-0	 kot-sa	kot-gu	 gaze-0	 gaze-sa	gaze-gu
 weil-0	 weil-sa	weil-gu	 balo-0	 balo-sa	balo-gu
 houf-0	 houf-sa	houf-gu	 cavu-0	 cavu-sa	cavu-gu

(b) Animacy-based gender system.

Table 1: Example training artificial lexica for sex-based (a) and animacy-based (b) suffixal gender systems in Experiments 1 and 2. During training, participants saw the zero-marked gender forms both in singular and plural, but the others (MASC or AN), they only saw in singular. During the testing phase, participants had to produce the forms for the held-out G2.PL (i.e., MASC.PL or AN.PL depending on the condition).

For each experiment, a total of 80 participants’ data was analysed (20 in each condition of our 2x2 factorial design).

Data Analysis

For each experiment, we use *R*’s *brms* (Bürkner, 2018) library as an interface to *Stan* (Carpenter et al., 2017) to fit a mixed-effects Bayesian binomial model predicting participants’ production of gender-closest order by condition (stem-G2-PL if suffixing or PL-G2-stem if prefixing). Our dependent variable is participants’ responses for each of the critical test trials (coded as 1 if gender-closest, and 0 if number-closest). As fixed effects, we include affix order (prefixing or suffixing) and gender system (animacy-based or sex-based), as well as their interaction. The categorical predictors are sum-coded, comparing each level to the grand mean. As random effects, we included intercepts for participants as well as stems. We set the same Student-*t* prior on all fixed effects and intercepts ($DF = 6, \mu = 0, \sigma = 1.5$); for random effects’ standard deviations, we set a half-Cauchy prior with a scale parameter 10.

Results

Based on our preregistered hypothesis, we predict that participants prefer to place gender morphology closer to the stem than number morphology in systems where both gender and number affixes are separate and both are placed before or after the stem. We also further explore whether participants are equally likely to produce gender-closest orders in systems where the same lexical stem can be inflected for different gender feature values (i.e., the sex-based system in our experiments) than in systems where lexical items can only be

inflected for a gender value (i.e., the animacy-based system in our experiments).

Figure 2 shows participants’ proportions of gender-closest orders in the critical trials. A visual inspection suggests that our results are consistent with the preregistered hypothesis: participants infer gender-closest orders (i.e., stem-G2-PL or PL-G2-stem), regardless of the affixes’ position or the type of gender system.

The model’s results for Exp 1 confirm our hypothesis: The intercept suggests very strong evidence that English speakers prefer to produce gender-closest as opposed to number-closest orders ($\hat{\beta} = 5.951$, 90%CI = [4.483, 7.810], $SE = 1.088$, $P(\hat{\beta} > 0) = 0.99$). The model further suggests that this preference is equally strong regardless of whether gender and number affixes are placed before or after the stem ($\hat{\beta} = 0.590$, 90%CI = [-0.165, 1.398], $SE = 0.478$, $P(\hat{\beta} > 0) = 0.90$) or whether they belong to the sex-based or the animacy-based system ($\hat{\beta} = 0.708$, 90%CI = [-0.237, 1.723], $SE = 0.612$, $P(\hat{\beta} > 0) = 0.89$). There is no evidence of an interaction between these two factors ($\hat{\beta} = 0.093$, 90%CI = [-0.708, 0.870], $SE = 0.048$, $P(\hat{\beta} > 0) = 0.58$). For English speakers, then, we do not find evidence of any difference across gender systems: Participants are equally likely to produce gender-closest orders for sex-based and for animacy-based systems.

The model’s results for Exp 2 also confirm our hypothesis: Italian speakers produce gender-closest orders notably above chance ($\hat{\beta} = 6.534$, 90%CI = [5.010, 8.505], $SE = 1.096$, $P(\hat{\beta} > 0) = 1$), and they do so equally regardless of whether gender and number morphology is prefixal or




Figure 2: Proportion of responses with gender-closest order in Exp 1 (A) and Exp 2 (B). Coloured dots represent participants' individual scores; black dots represent the model's predicted mean accuracy scores and the error bars represent the model's predicted 90% credible intervals. Dashed lines represent the chance level.

suffixal ($\hat{\beta} = 0.623$, $90\%CI = [-0.344, 1.653]$, $SE = 0.612$, $P(\hat{\beta} > 0) = 0.85$). However, unlike for English speakers, we find that Italian speakers are slightly more likely to produce gender-closest orders in the animacy-based system than in the sex-based system ($\hat{\beta} = 1.275$, $90\%CI = [0.280, 2.327]$, $SE = 0.630$, $P(\hat{\beta} > 0) = 0.98$), regardless of affix order ($\hat{\beta} = 0.129$, $90\%CI = [-0.845, 1.110]$, $SE = 0.593$, $P(\hat{\beta} > 0) = 0.59$).

Discussion

The experiments reported here tested whether language learners would infer that a novel gender morpheme should be ordered closer to the noun stem, with a novel number morpheme placed peripherally. Crucially, the experiments were designed so that no co-occurrence statistics would lead learners to prefer one order over the other (suggested as a mechanism for ordering preferences by e.g., Hahn et al., 2022). In other words, stems occurred equally frequently with the gender and number morpheme. Our results confirm a preference for gender closer to the stem than number, across two different linguistic populations with different gender systems in their first language. Participants infer the predicted order regardless of whether the morphemes were prefixal or suffixal, indicating that this preference is not about the sequential order of the morpheme—i.e., which comes first—but about which morpheme should be closer to the noun stem. We also found that the preference towards number-closest orders is consistent across sex-based variable gender, and animacy-based fixed gender. This suggests that the interpretation of gender across the board leads to a gender-first derivation, separate from number. This is consistent with theories that argue for a representation of gender together with (or closest to) the noun (Harris, 1991; Alexiadou, 2004; Mel'čuk, 2013; Kramer, 2015, 2016), but calls into question proposals that consider gender to always be bundled with number (e.g., Ritter, 1993; Carstens, 2003). This suggests that the pervasiveness of cumulative gender-number marking in affixes should be explained not by the abstract representation of these features, but by alternative historical mechanisms, such as phonological fusion.

However, we do find a slight difference in the likelihood of producing the reversed number-closest order in Exp 2 with Italian speakers: unlike English speakers, they are more likely to produce number-closest orders in variable sex-based systems. Unlike for animacy-based systems, a preference for number-closest is found for a small number of participants in the sex-based conditions. It is possible that gender morphemes are more likely to be interpreted as separate from the stem with stems that can derive related semantic entities such as the ones in our sex-based condition. This result is potentially in line with theories that posit that gender is derived with the noun stem in fixed systems and with number in variable systems (e.g., De Vincenzi, 1999). However, why do we not find this difference in English speakers is not clear. Given the strength of the overall preference, it is possible that our sample size is insufficient. However, it is also possible that it results from differences between these two linguistic populations. Unlike Italian speakers, English speakers do not have any experience with inflectional gender marking in their native language; it is thus possible that they are falling back on their knowledge of the placement of derivational morphology relative to inflectional morphology in sex-based systems. As is the norm with derivational morphology in general, derivational sex-based gender morphemes in English are attached closer to the noun stem than inflection (e.g., *lion-ess-(e)s*). If English speakers treat novel gender morphemes in our experiment as similar to English derivational morphology, then they may place it closer to stem in both systems tested here. Importantly, this is less likely to explain Italian speakers' behaviour. Italian has both fixed and variable gender (and even a derivational gender suffix, e.g., *leon-essale*), but inflectional gender and number morphology are always cumulative—i.e., there is no evidence for their relative order. While these results thus generally support the claim that the cognitive representation of gender and number impacts ordering preferences, in future work we will test additional populations to robustly assess the effect the type gender system (variable or fixed), and the impact of learners' prior linguistic knowledge.

Acknowledgements This work had received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 757643) and the NCCR Evolving Language, Swiss National Science Foundation (Agreement No 51NF40_180888).

Ethics The study was approved by the PPLS Research Ethics Committee at the University of Edinburgh (Ref No 423-2122/2) and the Ethics Committee of the School of Philosophy at the University of Zurich (Authorisation No 22.10.12).

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
- Bickel, B., Nichols, J., Zakharko, T., Witzlack-Makarevich, A., Hildebrandt, K., Rießler, M., ... Lowe, J. B. (2022). *The autotyp database (v1.1.0)*. Retrieved from <https://doi.org/10.5281/zenodo.6793367> doi: 10.5281/zenodo.6793367
- Bybee, J. L. (1985). *Morphology: A study of the relation between meaning and form*. Philadelphia, PA: John Benjamins.
- 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
- Carminati, M. N. (2005). Processing reflexes of the feature hierarchy (person_i number_i gender) and implications for linguistic theory. *Lingua*, 115(3), 259–285. doi: 10.1016/j.lingua.2003.10.006
- 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(1). 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: 10.1162/002438903322247533
- 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. In M. S. Dryer & M. Haspelmath (Eds.), *The world atlas of language structures online*. Leipzig: Max Planck Institute for Evolutionary Anthropology. Retrieved from <https://wals.info/chapter/31>
- De Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior research methods*, 47(1), 1–12. doi: 10.3758/s13428-014-0458-y
- De Vincenzi, M. (1999). Differences between the morphology of gender and number: Evidence from establishing coreferences. *Journal of Psycholinguistic Research*, 28, 537–553. doi: 10.1023/A:1023272511427
- Gibson, E., Piantadosi, S. T., Brink, K., Bergen, L., Lim, E., & Saxe, R. (2013). A noisy-channel account of crosslinguistic word-order variation. *Psychological science*, 24(7), 1079–88. doi: 10.1177/09567976124637
- Greenberg, J. (1963). Some universals of grammar with particular reference to the order of meaningful elements. In J. Greenberg (Ed.), *Universals of language* (p. 73-113). Cambridge, MA: MIT Press.
- Hahn, M., Degen, J., & Futrell, R. (2021). Modeling word and morpheme order in natural language as an efficient tradeoff of memory and surprisal. *Psychological Review*, 128(4). doi: 10.1037/rev0000269
- 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
- Harris, J. W. (1991). The exponence of gender in Spanish. *Linguistic inquiry*, 22(1), 27–62. Retrieved from <https://www.jstor.org/stable/4178707>
- Hawkins, J. A., & Cutler, A. (1988). Psycholinguistic factors in morphological asymmetry. In J. A. Hawkins (Ed.), *Explaining language universals* (pp. 280–317). Oxford: Blackwell.
- Hay, J. (2001). Lexical frequency in morphology: Is everything relative? *Linguistics*, 39(6), 1041–1070. doi: 10.1515/ling.2001.041
- 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* (Vol. 58). Oxford University Press.
- Kramer, R. (2016). The location of gender features in the syntax. *Language and linguistics Compass*, 10(11), 661–677. doi: 10.1111/lnc3.12226
- 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* (Vol. 2, pp. 163–176). Amherst, MA: University of Massachusetts GLSA. (Preprint doi: 10.31234/osf.io/5s2r8)
- 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
- Mel’čuk, I. A. (2013). The notion of inflection and the expression of nominal gender in Spanish. *Studies in Language. International Journal sponsored by the Foundation “Foundations of Language”*, 37(4), 736–763. doi: 10.1075/sl.37.4.02mel
- Picallo, M. C. (1991). Nominals and nominalizations in catalan. *Probus*, 3(3), 279–316. doi: 10.1515/prbs.1991.3.3

- 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).
- Ritter, E. (1993). Where's gender? *Linguistic inquiry*, 24(4), 795–803. Retrieved from <https://www.jstor.org/stable/4178843>
- 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