Singling out singular: Exploring cognitive biases underlying probable category partitions in nominal number systems

Carmen Saldana (carmen.saldana@ub.edu)[1,2], Emma Angela Montecchiari[2,3], Jennifer Culbertson[4] & Mora Maldonado[2]

[1]Department of Catalan Philology and General Linguistics, Universitat de Barcelona
[2]Department of Translation and Language Sciences, Universitat Pompeu Fabra
[3]Center for Mind/Brain Sciences, Università degli Studi di Trento
[4]Center for Language Evolution, The University of Edinburgh

Abstract

Grammatical number is one of the most widespread morphological categories in the world's languages. When only a single number distinction is expressed, it is almost always between singular and non-singular (plural), and even in systems with additional values such as dual, non-singular categories are frequently collapsed. The collapse of singular and dual categories is less frequent but attested, whereas groupings that merge singular with plural are virtually unattested. These patterns reveal strong cross-linguistic asymmetries in how number systems are partitioned. Linguistic theories have proposed that such regularities reflect underlying feature primitives, such as SINGULAR and MINIMAL (Harbour, 2014; Noyer, 1992; Silverstein, 1976). The former explains why dual and plural constitute a natural non-singularity category, and the latter explains why singular and dual do so, as both are the cardinally smallest singular and non-singular values. However, little is known about whether these formal features correspond to learners' cognitive representations. In this paper, we test this proposal using two artificial language learning experiments that manipulate alternative number partitions. Our results show that learners systematically prefer partitions consistent with the proposed feature structure, with strongest preferences for SINGULAR-based contrasts, followed by MINIMAL-based distinctions. These findings provide initial evidence for a cognitively principled explanation of the typology of partitions of grammatical number systems.

Keywords: number; morphology; features; typology; artificial language learning

Introduction

Languages vary greatly in the grammatical categories they encode, but some are far more pervasive cross-linguistically than others, and grammatical number is among the most widespread (Acquaviva & Daniel, 2022; Bickel & Nichols, 2007; Corbett, 2000; Skirgård et al., 2023). Most documented languages mark at least one number distinction in pronominal and/or nominal systems, and when only a single distinction is encoded, it is between singular (SG) versus plural (PL).

Most languages productively mark plurality on nouns or noun phrases. In Grambank (Skirgård et al., 2023), fewer than one third of documented systems show no nominal plural marking (31%, $n = 572$ of 1866 languages with complete information), and in roughly 60% of all nominal systems documented the only contrast is SG/PL. Only a minority of languages mark more than one non-singular category (11% of languages in Grambank); in these cases, most systems (79%) contrast dual (DU) and PL uniquely. And even in these three value systems (SG/DU/PL), syncretism among non-singular categories (i.e., the use of the same form for different non-singular categories, resulting in DU = PL in a *subset* of the linguistics system) is very frequent (Baerman et al., 2005). Taken together, these patterns reveal a strong cross-linguistic tendency to uniquely distinguish SG from PL, and consolidate non-singular values into a single, natural category.

We have limited cross-linguistic evidence for syncretism between singular and non-singular values that do not result in full number neutralisation. However, what evidence there is indicates that while syncretism between SG and DU does occur, syncretism between SG and PL in systems that distinguish more than one non-singular category (e.g., SG/DU/PL) is virtually unattested (Baerman et al., 2005; Lee, 2020). This asymmetry suggests that SG and DU may also form a more natural grouping than SG and PL.

Typological evidence of this sort supports the assumption that number feature values such as SG or PL are not *sui generis*, but rather composites of more fundamental feature primitives (Hale, 1973; Harbour, 2011; Noyer, 1992; Silverstein, 1976). If one assumes features like SG, DU and PL without further decomposition, there is no explanation for why categories that collapse DU and PL, or SG and DU are more natural and common cross-linguistically. However, if we introduce a more fundamental feature such as MINIMAL, which is positively valued when it refers to the cardinally smallest set within either singular and non-singular categories, we can decompose DU into {-SINGULAR, +MINIMAL} as it is non-singular but it refers to the cardinally smallest possible non-singular category. We then decompose SG into {+SINGULAR, +MINIMAL}, and PL into {-SINGULAR, -MINIMAL}. These feature decompositions would suffice to explain why DU=PL are a natural class defined by {-SINGULAR}, as well as why SG=DU are a natural class defined by {+MINIMAL} (Harbour, 2011; Noyer, 1992; Silverstein, 1976). Moreover, some proposals suggest that these features apply recursively (Harbour, 2011, 2014). First, SINGULAR partitions the space into singular and non-singular, then MINIMAL applies to the non-singular partition yielding the cardinally smallest group, that is, the dyad referred to by DU: (+MINIMAL(-SINGULAR)). If we apply MINIMAL further recursively (Harbour, 2011), we get to the triad referred to by the rarer trial (TRI) category: (+MINIMAL(-MINIMAL(-SINGULAR))).

While these theories of feature decomposition can successfully account for the cross-linguistic structure of gram-

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

matical number systems involving SG, PL and even DU and TRI categories, little is known about whether these formal features capture speakers' cognitive representations, which would yield the aforementioned natural classes cognitively—and not merely formally—natural. If we can show that these reflect speakers' representations and impact language learning, we can provide evidence for a cognitively-principled explanation of the typology of partitions of the grammatical number space (Culbertson, 2012; Gentner & Bowerman, 2009; Saldana et al., 2022).

One way to probe these biases is to assess differences in speakers' preferences across the proposed more or less natural partitions of the number space. Following the aforementioned feature structure, we would expect SG versus PL partitions (defined by a ±SINGULAR) to be most readily inferred and learned, followed by {SG,DU} versus PL partitions; unnatural partitions such as {SG,PL} vs DU should be most dispreferred.

Following previous work exploring the relationship between learning and morphological systems (e.g., Lee, 2020; Maldonado & Culbertson, 2021; Saldana et al., 2022, 2025), we use an artificial language learning paradigm to test the predictions outlined above for nominal number marking systems. In Experiment 1, we compare adult learners' ease of acquisition of systems with homophony patterns which differ in their formal naturalness and typological frequency: SG VS. DU=PL, SG=DU VS. PL and SG=PL VS. DU.

Experiment 1

Materials and Methods

Participants (Spanish L1 speakers) are trained and tested on a miniature language expressing grammatical number distinctions, i.e., a set of nouns with number marking. We manipulate between participants the pattern of homophony present in the system, creating conditions instantiating three different partitions that vary in their formal naturalness and typological frequency. We analyse the differences in learnability across these partitions to evaluate their cognitive naturalness. The preregistered hypotheses, design and analysis plan are accessible at https://osf.io/cy5vh/.

Experimental conditions. The artificial languages in each condition model a grammatical number system with complete (whole-word) homophony between two of three values—SG (i.e., a single entity), DU (i.e., a group of two entities), and PL (a group of more than two entities)[1]—resulting in the neutralisation of one number distinction. Across the three conditions, we manipulate homophony patterns (i.e., between SG and DU, between DU and PL, or between SG and PL). Participants are taught number systems with only two affixes, referred to in the text as A or B. One of the affixes corresponds to a single number value, the other affix corresponds to the two remaining values. The three conditions are thus the following: SG VS. DU=PL (ABB hereafter, i.e., a +SINGULAR/-SINGULAR partition), SG=DU VS. PL (AAB, i.e., a +MINIMAL/-MINIMAL partition), and SG=PL VS. DU (ABA hereafter, an unnatural partition). Participants are randomly assigned to one of the three conditions. Number affixes vary between participants, and can be either prefixes or suffixes; participants are randomly assigned an affix order at the beginning of the experiment. This allows us to further assess whether the strength of the learning preference depends on the familiarity of affix position in the expression of grammatical number (i.e., stronger for suffixal systems, which mirror participants' L1 systems).

Artificial language. The artificial language lexicon comprises six disyllabic stems and two monosyllabic CV affixes, randomly selected from the set {*ne, pu, da, si, to*}. Stems refer to entities from three animacy classes (human, non-human animate, and inanimate), with two stems from each class assigned to each participant. Stems are drawn from a set including *kuki* ('cook'), *doco* ('doctor'), *juve* ('teenager'), *kria* ('toddler'), and *ancio* ('elder') for human entities; *mise* ('cat'), *cane* ('dog'), *kaba* ('goat'), *beehbe* ('sheep') and *purco* ('pig') for non-human animates; and *burso* ('backpack'), *vala* ('ball'), *xarre* ('mug'), *botil* ('bottle') and *lepi* ('pencil') for inanimates. Lexical items with affixes are used to refer to simple images of the entities above in different cardinalities.

Experimental procedure. Participants are asked to learn the mapping between meanings and forms via feedback-learning. On each trial, participants complete a two-alternative forced-choice task: they are presented with an image of a specific cardinality and two forms which share the correct stem but differ in the affix forms attached to them (i.e., one with affix A and the other with affix B). Participants must select the correct corresponding form. They receive feedback after their selection. Initially, participants have to guess the mappings, and they learn the correct form-meaning mappings over time via the feedback provided. Participants receive a bonus of £0.01 for each correct response.

Participants are exposed to 9 blocks of six trials each, with two trials for each number value SG, DU, PL (a total of 54 trials). Each block contains trials with two different referent entities. These 9 blocks are organised into 3 meta-blocks of 18 trials, each including all six referent entities, and each appearing with the different number values the same number of times (i.e., 6 entities x 3 number values). Thus, each entity is presented 9 times across the experiment, with each number value appearing three times.

In all conditions, the first block is designed as an ABX-type extrapolation task. Participants first complete four trials involving form–meaning mappings for only two number values, one mapped to affix A and the other to affix B. In the final two trials, they are asked to select the form corresponding to the remaining, previously unseen number value (see Table 1 for the possible extrapolation types). The fifth trial of this first block thus constitutes the critical X trial, allowing us to examine participants' initial generalisation (i.e., their guess) as to whether the affix A or B map most naturally to the novel cardinality. In other words, these trials measure participants' inferences about the structure of the number meaning space.

[1]The cardinality of plural referents (i.e., > 2) varies between participants, and can be either 3 only, 4 only, 5 only, 3-4 or 4-5.

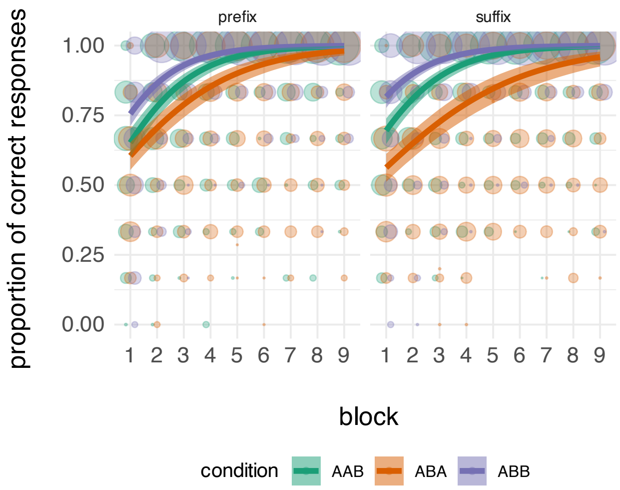

Figure 1: Accuracy by testing block for each of the three conditions in Experiment 1, across prefixal and suffixal number marking systems. Shaded dots represent participants' individual scores, and larger dots represent more individuals; thick lines represent the model's predicted mean accuracy conditioned on experimental condition and block. The shaded area shows the 90% credible intervals.

Note that for each condition, there were two possible extrapolation types (e.g., in the ABB condition, participants could either first be trained on SG/DU and then extrapolation to PL, or first be trained on SG/PL and then extrapolation to DU); the extrapolation type was randomly assigned for each participant.

Participants Participants (N = 507) are adult Spanish speakers recruited via Prolific (filters: Spanish as a first and primary language) for a study lasting a median of 7-8 minutes. They are compensated a base rate of £1.20, and can get up to £0.54 in bonuses according to performance (£0.01 per correct trial). We exclude participants whose translations of the number markers in the post-experimental survey do not mention number-related meanings, and those who report they have taken notes during the experiment. After exclusions (N = 33) we have 157, 157, and 160 participants in conditions ABB, AAB and ABA, respectively (total of N = 474). They are evenly distributed across prefix and suffix orders.[2]

Results

Learnability Recall that we predicted a preference for ABB partitions (e.g., a +SG/-SG partition), followed by a preference for AAB partitions (e.g., a +MIN/-MIN partition); we expect ABA partitions to be the least preferred. To test whether our results match a learnability gradient, ABB > AAB > ABA, we fitted a Bayesian hierarchical binomial regression model (with a logit link) using `brms` (Bürkner, 2017, 2018) as an interface to `Stan` (Carpenter et al., 2017). Our dependent variable is participants' response at each trial, either correct or incorrect. As fixed effects we include block (centered), condition (AAB [reference level], ABB and ABA) and system type (suffixal [reference level], and prefixal) with interaction terms. We also include random intercepts by participant and by the cardinality of the plural referent, as well as by-participant random slopes for the effect of block.

Results are in line with our predictions (see Figure 1). We find that overall accuracy at the intercept (block 5) for ABA is lower than for AAB ($\hat{\beta} = -1.628$, 90%CI = $[-2.124, -1.137]$, $P(\hat{\beta} < 0) = 1$), while accuracy for ABB is higher than for AAB ($\hat{\beta} = 0.926$, 90%CI = $[0.408, 1.455]$, $P(\hat{\beta} > 0) = 0.998$). As expected, accuracy increases over blocks ($\hat{\beta} = 0.627$, 90%CI = $[0.541, 0.713]$, $P(\hat{\beta} > 0) = 1$), though the learning rate is slower for ABA compared to AAB ($\hat{\beta} = -0.264$, 90%CI = $[-0.377, -0.154]$, $P(\hat{\beta} > 0) < 0.001$), while the slope for ABB is comparable to AAB ($\hat{\beta} = 0.064$, 90%CI = $[-0.060, 0.188]$, $P(\hat{\beta} > 0) = 0.8$). We do not find any evidence for an effect of affix order in either AAB or ABB.

Extrapolation Aside from examining the learnability of different partitions of the number space, we also tested partici-

Extrapolation type A/B → **X**	Alternative partitions	Predicted partition
SG/PL → DU	A**B**B vs AA**B**	A**BB** (i.e., DU = PL)
SG/DU → PL	AB**B** vs AB**A**	A**BB** (i.e., DU = PL)
DU/PL → SG	**A**AB vs **A**BA	**AA**B (i.e., SG = DU)

Table 1: Extrapolation types in Experiment 1. In the first block of trials, participants are trained on two number values, corresponding to A and B affixes, and then must extrapolate A or B to the third value. For each type, alternative partitions at the time of extrapolation are shown (extrapolated form in bold) along with the predicted responses.

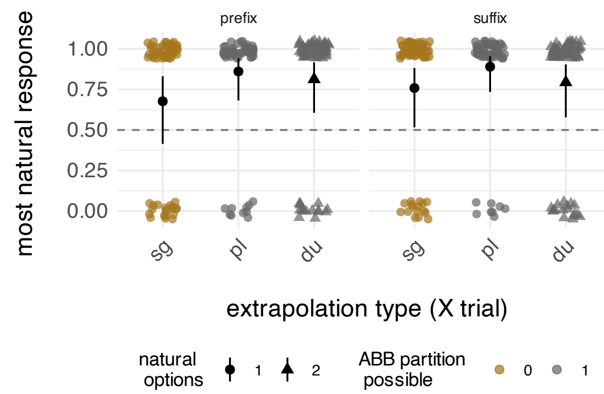

Figure 2: Participants' selection of the most natural response by extrapolation type in Experiment 1 (i.e., extrapolation to SG, DU or PL, see Table 1). Shaded dots and triangles represent participants' individual responses; black solid dots represent the model's predicted selection means. The error bars are the 90% CIs. Dot colour indicates whether an ABB partition (+SG/-SG) is available (1) or not (0) for that extrapolation type; shape indicates whether the two alternative partitions at the point of extrapolation are both natural (2) or whether only one is natural (1). The horizontal dotted line indicates chance.

[2]They are also evenly distributed across the different cardinalities mapped to the plural value (see Footnote 1); for each of the 10 different parameter combinations per condition (2 affix orders, 5 different plural referents), we have between 15 and 18 participants.

pants' initial inferences before they were exposed to the full number marking system. We predicted differences in the likelihood of extrapolation by partition type. Responses conforming to an ABB (+SG/-SG) partition should be most likely. In this case, it is also the closest match to participants' L1 number marking system. When this partition is not available, the next most likely response is predicted to be a AAB (+MIN/-MIN). Unnatural partitions ABA are predicted to be the least likely. To test these predictions, we fitted a Bayesian hierarchical binomial regression model predicting whether participants' response in the first extrapolation trial matched our predicted partition, by extrapolation type (see Table 1) and affix order and their interaction. We also include intercepts by plural cardinality, condition and their interaction.

The models' results (see Figure 2) suggest that when extrapolating to SG—with AAB and ABA as alternative partitions—participants are most likely to select the most natural pattern (i.e., AAB in this case; $\hat{\beta}$ = 1.104, 90%CI = [0.065, 2.007], $P(\hat{\beta} > 0) = 0.957$). Extrapolation to DU—with ABB and AAB as alternative partitions—yields comparably high mean estimates of most natural response (i.e., ABB in this case; $\hat{\beta}$ = 0.221, 90%CI = [−0.472, 0.919], $P(\hat{\beta} > 0) = 0.696$), and in extrapolations to PL—with ABB and ABA as alternative partitions—participants are even more likely to choose the most natural response (i.e., ABB; $\hat{\beta}$ = 0.966, 90%CI = [0.240, 1.708], $P(\hat{\beta} > 0) = 0.986$). We again find no evidence for an effect of affix order (prefix vs. suffix). Taken together, these results indicate that participants tend to infer natural partitions, and when possible, the SG/¬SG partition. Moreover, the tendency to infer an ABB (+SG/-SG) partition increases when the alternative is the unnatural partition ABA.

Interim Discussion

In this paper we asked whether the typological dominance of certain partitions of the grammatical number systems, which are deemed most natural formally, can be linked to cognitive biases active during language learning. In Experiment 1, participants learned grammatical number systems referring to cardinalities within the range {1,2,>2}. We found the predicted learnability gradient ABB (+SG/-SG) > AAB (+MIN/-MIN) > ABA (unnatural). We also found that, in the absence of evidence of the full morphological system, participants always infer a natural partition. When an +SG/-SG partition is available, they infer that, even if in competition with +MIN/-MIN; when it is not and the only natural partition available is +MIN/-MIN, they infer the latter. These preferences are robust and across prefixal and suffixal number marking systems. These results confirm that what are defined as natural classes formally with a proposed feature structure {SINGULAR,MINIMAL} are also natural categories for learners acquiring a grammatical number system.

Results also suggest that a +SG/-SG is most preferred, which is also in line with the proposed recursive application of {SINGULAR} and {MINIMAL} features, in that order (Harbour, 2011); however, +SG/-SG is the number marking system in their first language—i.e., Spanish, with a SG/PL distinction—and could be due to structural transfer effects alone. Further work should seek to extend these results to populations of learners whose grammatical number systems differ. Stronger evidence for the privileged status of +SG/-SG partitions would come from participants whose first language features DU (e.g., Slovene or Gulf Arabic), a number marking system with a +MIN/-MIN partition (i.e., SG=DU/PL, a productive syncretism in Central Pame), or a lack of systematic number marking (e.g., Japanese or Mandarin Chinese).

Another way of assessing the privileged status of +SG/-SG within the same population, however, is to compare the preference for natural partitions in novel cardinality ranges where either individual entities SG or dyads DU are present as the cardinally smallest value. For example, we can compare the observed learnability preference towards ABB partitions across cardinality ranges such as {1,3,5} and {2,4,6}. If an +SG/-SG is more readily available to participants than +MIN/-MIN, we would expect a higher preference for ABB patterns with triplets like {1,3,5}—where the pattern is defined by +SG/-SG as in Experiment 1—than with triplets like {2,4,6}—where it is defined by +MIN/-MIN. We would also expect a lower preference for ABB with cardinality triplets whereby the cardinally smallest value is a triad, for example, {3,5,7}. If we apply MINIMAL further recursively (Harbour, 2011, 2014), we get to the triad referred to by TRI as follows: (+MINIMAL(-MINIMAL(-SINGULAR))). We could then expect the following preference gradient for ABB patterns across triplets: {1,3,5} > {2,4,6} > {3,5,7}. With the assumed recursive MINIMAL feature, the learnability gradient ABB > AAB > ABA would nevertheless be predicted across all triplets: contiguous values for each triplet could form a natural class—predicting ABB|AAB > ABA—and those involving less recursive applications would be preferred—ABB > AAB. To test these hypotheses, in Experiment 2 we examine learnability partitions in which the minimal cardinality varies between 1, 2, or 3.

Experiment 2

Materials and Methods

We use the same exact design and trial structure as in Experiment 1 to test the learnability of ABB, AAB, and ABA partitions with the following three cardinality triplets: {1,3,5}, {2,4,6}, and {3,5,7}. Note that cardinalities within triplets are either all even or all odd, and the intervals between cardinalities are constant. Unlike in Experiment 1, here the lowest cardinality varies across triplets. Therefore, instead of referring to specific number values, here ABB, AAB, and ABA refer to the sequence of LOW, MID and HIGH cardinalities in the triplet, in that order. We test only suffixal systems. Participants are randomly assigned to one of the three conditions (i.e., ABB, AAB and ABA), with one of the three cardinality triplets, for a total of nine experimental groups.

Participants Participants (N = 221) are adult Spanish speakers recruited via Prolific (filters: Spanish as first and primary language) for a study lasting a median of 7-8 minutes. They

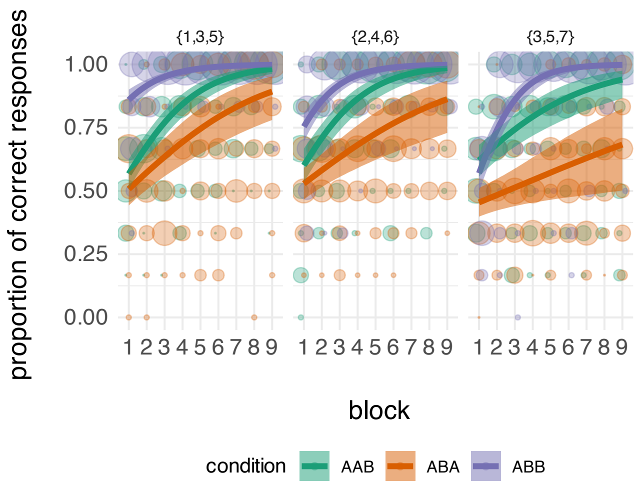

Figure 3: Accuracy by testing block for each of the three conditions and cardinality triplets in Experiment 2.

are compensated a base rate of £1.20, and can get up to £0.54 in bonuses according to performance (£0.01 per correct trial). After exclusions as per the previous experiment (N = 30), we analyse data from a total of 191 participants—i.e., 20-24 participants for each of the nine conditions x cardinality triplets.

Results

Learnability We fitted a Bayesian hierarchical binomial regression model predicting trial-level accuracy from block, condition (ABB/AAB/ABA), and cardinality triplet ({1,3,5}/{2,4,6}/{3,5,7}) and their interactions, with by-participant intercepts and by-participant slopes for BLOCK. Figure 3 shows raw accuracy together with model-predicted means and 90% CIs. Overall performance in the baseline (AAB, {1,3,5}) was well above chance at the intercept (block 5; $\hat{\beta}$ = 2.153, 90% CI [1.736, 2.582], $P(\hat{\beta} > 0) = 1$). Relative to AAB, accuracy was higher in the ABB condition ($\hat{\beta}$ = 1.986, 90% CI [1.285, 2.708], $P(\hat{\beta} > 0) = 1$) and lower in ABA ($\hat{\beta} = -1.073$, 90% CI [−1.701, −0.459], $P(\hat{\beta} < 0) = 0.998$). Accuracy increased across blocks ($\hat{\beta}$ = 0.470, 90% CI [0.371, 0.572], $P(\hat{\beta} > 0) = 1$), with slower learning in ABA relative to AAB ($\hat{\beta} = -0.208$, 90% CI [−0.347, −0.069], $P(\hat{\beta} < 0) = 0.995$), and a comparable learning rate in ABB. We found no evidence for a condition by triplet interaction, indicating a robust learnability ordering ABB > AAB > ABA across numerical ranges.

Although the overall condition effect was stable across triplets, early performance (block 1) revealed triplet-specific differences. For the {1,3,5} and {2,4,6} triplets, there was strong posterior support for an early advantage of ABB over AAB (respectively, $P(\text{ABB} > \text{AAB}) = 1$ and $P(\text{ABB} > \text{AAB}) = 0.997$). By contrast, for the {3,5,7} triplet there was no posterior support for an early preference of ABB over AAB ($P(\text{ABB} > \text{AAB}) = 0.19$). In all three triplets, ABA was disfavoured relative to AAB.

Comparisons across triplets within each condition at block 1 showed that triplet effects were particularly relevant in ABB.

Extrapolation type A/B → **X**	Alternative partitions	Predicted partition
LOW/HIGH → MID	A**B**B vs AA**B**	A**B**B
LOW/MID → HIGH	AB**B** vs AB**A**	AB**B**
MID/HIGH → LOW	**A**AB vs **A**BA	**A**AB

Table 2: Extrapolation types and predictions in Experiment 2.

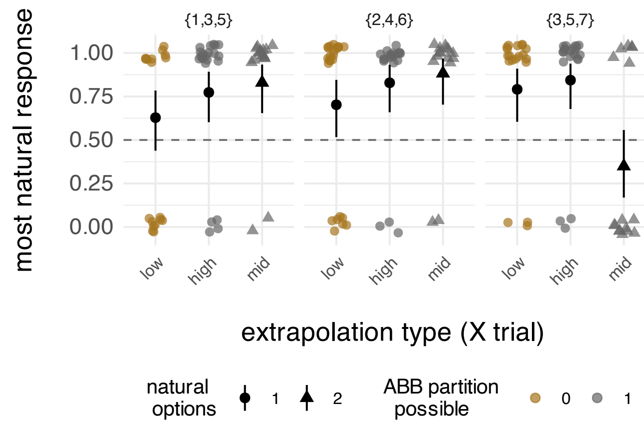

Figure 4: Participants' selection of the most natural response by extrapolation type in Experiment 2 (i.e., extrapolation to LOW, MID or HIGH, see Table 2). Colour indicates whether an ABB partition is available, and shape indicates the number of natural partitions available at the point of extrapolation.

Posterior support indicated a graded decline in ABB accuracy as the minimal cardinality increased ({1,3,5} > {2,4,6} > {3,5,7}; $P(\{1,3,5\} > \{2,4,6\}) = 0.991$, $P(\{1,3,5\} > \{3,5,7\}) = 1$, $P(\{2,4,6\} > \{3,5,7\}) = 0.999$). By contrast, AAB showed no comparable sensitivity to triplet, and ABA only weak differences. Together, these results indicate that while the natural > unnatural preference is comparable across numerical ranges, the strength of a preference ABB > AAB decreases as the minimal cardinality in the triplet increases. In other words, the bias is strongest when the system can be defined by {singular}, followed by {minimal}, and weakest when the partition needs to be defined by further recursive applications of {minimal}.

Extrapolation We fitted a Bayesian binomial regression model to assess what pattern of number marking participants inferred in the first extrapolation trial (see Figure 4). As in Experiment 1, our dependent variable is whether participants selected the predicted partition. We predict the ABB partition to be the most likely choice. When this partition is not available, the next most likely choice is the AAB partition. Unnatural ABA partitions should be the least likely to be inferred (see Table 2).

We find that at the baseline levels (extrapolation to LOW with the {1,3,5} triplet), participants are somewhat more likely than chance to select the predicted response ($\hat{\beta}$ = 0.525, 90%CI = [−0.248, 1.289], $P(\hat{\beta} > 0) = 0.878$). In this case, ABB is not available, and the preferred pattern is

AAB (rather than ABA). For the same triplet {1,3,5}, extrapolations to MID and HIGH show a higher likelihood of selecting the predicted pattern (MID: $\hat{\beta} = 1.074$, 90%CI = [0.056, 2.176], $P(\hat{\beta} > 0) = 0.957$); HIGH: $\hat{\beta} = 0.709$, 90%CI = [−0.205, 1.635], $P(\hat{\beta} > 0) = 0.899$). In both these cases the ABB pattern is available. The selection of predicted responses in extrapolations to LOW is comparable for {2,4,6} ($\hat{\beta} = 0.339$, 90%CI = [−0.496, 1.201], $P(\hat{\beta} > 0) = 0.743$) but higher for {3,5,7} ($\hat{\beta} = 0.816$, 90%CI = [−0.127, 1.817], $P(\hat{\beta} > 0) = 0.921$). The selection of the predicted response in extrapolation to MID for the triplet {3,5,7} is notably lower ($\hat{\beta} = -3.061$, 90%CI = [−4.601, −1.620], $P(\hat{\beta} < 0) = 1$).

Altogether these results suggest that participants infer the predicted partitions across extrapolation types for triplets {1,3,5} and {2,4,6}. Participants prefer natural partitions (ABB and AAB) over unnatural partitions (ABA), and ABB over AAB partitions. However, results differ for the triplet {3,5,7}: participants prefer natural over unnatural partitions, but they are not more likely to infer ABB over AAB partitions.

General Discussion

This paper asked whether the typology of nominal number systems can be linked to learners' preferences for natural partitions of the number space, that is, partitions that correspond to natural classes definable in terms of feature primitives such as SINGULAR and MINIMAL, and its recursive application (Harbour, 2011, 2014; Noyer, 1992; Silverstein, 1976). Across two artificial language learning experiments, we find converging evidence for both a bias towards the hypothesised natural partitions and a more specific preference for those partitions defined by SINGULAR (+SG/-SG) over those defined by MINIMAL (+MIN/-MIN) and their recursive application—i.e., contrasting (+MIN(-MIN)) vs. (-MIN(-MIN)).

In Experiment 1, we tested the learnability gradient ABB > AAB > ABA in {1,2,>2} cardinality ranges. Here, proposed unnatural partitions ABA were dispreferred relative to natural ABB or AAB partitions both during learning and extrapolation. This preference was robust across prefixal and suffixal marking, suggesting that it does not depend on surface morphological order. Experiment 2 further shows that the same advantage holds across different cardinality ranges: for all three triplets ({1,3,5}, {2,4,6}, {3,5,7}), participants learned natural partitions better than unnatural alternatives. Together, these results support the hypothesis that the proposed feature primitives SINGULAR and MINIMAL and the resulting natural classes from their application (i.e., DU=PL or SG=DU) could constitute a cognitively privileged way of carving up the number space.

Beyond the general preference for natural partitions, we tested whether learners favour partitions solely defined by SINGULAR over those partitions solely defined by MINIMAL. In Experiment 1, this prediction aligns with the well-attested +SG/-SG preference: when +SG/-SG partition competed with +MIN/-MIN partition, participants preferred the former. Crucially, Experiment 2 further shows that the preference towards ABB partitions generalises across different numerical ranges—but only under certain conditions. For the {1,3,5} and {2,4,6} triplets, ABB partitions were both initially inferred and learned more readily than AAB partitions. In contrast, for the {3,5,7} triplet, participants continued to favour natural over unnatural partitions, but the preference for ABB partitions was absent in initial inference and in the earliest stages of learning. This provides evidence in support of a more readily available +SG/-SG partition to participants, followed by +MIN/-MIN, followed by (+MIN(-MIN))/(-MIN(-MIN)).

These results raise two related open questions about the nature of the primitives that underpin the proposed natural partitions. On the one hand, our results cannot shed light on whether the observed learning biases are specific to our participants' prior linguistic knowledge or whether they tap into more general representations which should be observed regardless of participants' linguistic systems. As discussed in our Interim Discussion, if these biases turn out to be robustly found across populations of learners whose grammatical number systems differ from the prototypical SG/PL partition, our hypothesis would be substantially strengthened. On the other hand, our results can neither confirm whether the preferences at play reflect linguistic representations of number or whether they mirror domain-general conceptual primitives which could apply equally to partitions of the linguistic and non-linguistic number space. It is possible that the results we obtain can be explained by domain general biases favouring contiguous over discontinuous partitions of meaning spaces, which have been shown to be at play in concept and category learning (Bruner et al., 1956; Goodman et al., 2008; Shepard et al., 1961, e.g.,), as well as in word learning (e.g., Carr et al., 2020; Dautriche et al., 2016; Landau & Shipley, 2001; Xu & Tenenbaum, 2007). This could explain the preference for our hypothesised natural over unnatural partitions. Moreover, the gradable preference for ABB partitions in Experiment 2 could be due to biases related to numerical cognition. One possibility is that small numerosities are differentially salient (e.g., 1 > 2 > 3), such that isolating the minimal value becomes less favoured as the minimal cardinality increases. A second, non-exclusive possibility is that categories not referring to exact numerosities (i.e., PL) which are lower bounded (e.g., ≥ 2) become less conceptually natural as their lower bound increases: a category such as "≥ 5"—which would correspond to the non-TRI category in {3,5,7} triplets—may be harder to treat as a coherent default than "≥ 3" or "≥ 4", as the lower bound approaches the limits of humans' subitising range (usually ≤ 4; Kaufman et al., 1949; Leibovich-Raveh et al., 2018; Trick & Pylyshyn, 1994). Addressing this confound will require extending the design to further numerical ranges.

Taken together, our findings provide initial evidence for a cognitively-principled explanation of the typology of partitions of grammatical number systems. Further work, however, will need to test whether the observed biases are shared across diverse linguistic populations, and whether they are specific to the linguistic domain.

Acknowledgements

This research has been supported by the Beatriu de Pinós programme (grant agreement No. 2022 BP 00101, held by Carmen Saldana) co-funded by the *Agència de Gestió d'Ajuts Universitaris i de Recerca* (AGAUR, Generalitat de Catalunya) and the European Union's Horizon 2020 research and innovation MSCA programme under agreement No 801370.

Ethics

The study was approved by the Institutional Committee for Ethical Review of Projects (CIREP) at Universitat Pompeu Fabra (Reference Nr. 345, PI: Carmen Saldana).

References

Acquaviva, P., & Daniel, M. (Eds.). (2022). *A comparative handbook*. De Gruyter Mouton. https://doi.org/doi:10.1515/9783110622713

Baerman, M., Brown, D., & Corbett, G. G. (2005). *The syntax-morphology interface: A study of syncretism* (Vol. 109). Cambridge University Press.

Bickel, B., & Nichols, J. (2007). Inflectional morphology [2nd edition]. In T. Shopen (Ed.), *Language typology and syntactic description*. Cambridge University Press.

Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A study of thinking*. John Wiley; sons, Incorporated.

Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, *80*(1), 1–28. https://doi.org/http://dx.doi.org/10.18637/jss.v080.i01

Bürkner, P.-C. (2018). Advanced Bayesian multilevel modeling with the R package brms. *The R Journal*, *10*(1), 395–411. https://doi.org/http://dx.doi.org/10.32614/RJ-2018-017

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, *76*(1).

Carr, J. W., Smith, K., Culbertson, J., & Kirby, S. (2020). Simplicity and informativeness in semantic category systems. *Cognition*, *202*, 104289. https://doi.org/http://dx.doi.org/10.1016/j.cognition.2020.104289

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

Culbertson, J. (2012). Typological universals as reflections of biased learning: Evidence from artificial language learning. *Language and Linguistics Compass*, *6*(5), 310–329.

Dautriche, I., Chemla, E., & Christophe, A. (2016). Word learning: Homophony and the distribution of learning exemplars. *Language Learning and Development*, *12*(3), 231–251. https://doi.org/http://dx.doi.org/10.1080/15475441.2015.1127163

Gentner, D., & Bowerman, M. (2009). Why some spatial semantic categories are harder to learn than others: The typological prevalence hypothesis. In *Crosslinguistic approaches to the psychology of language: Research in the tradition of dan isaac slobin* (pp. 465–480). Psychology Press.

Goodman, N. D., Tenenbaum, J. B., Feldman, J., & Griffiths, T. L. (2008). A rational analysis of rule-based concept learning. *Cognitive science*, *32*(1), 108–154.

Hale, K. L. (1973). Person marking in walbiri. In S. R. Anderso & P. Kiparsky (Eds.), *A festschrift for morris halle* (pp. 308–344). New York: Holt, Rinehart, Winston.

Harbour, D. (2011). Descriptive and explanatory markedness. *Morphology*, *21*(2), 223–245.

Harbour, D. (2014). Paucity, abundance, and the theory of number. *Language*, *90*(1), 185–229.

Kaufman, E. L., Lord, M. W., Reese, T. W., & Volkmann, J. (1949). The discrimination of visual number. *The American journal of psychology*, *62*(4), 498–525.

Landau, B., & Shipley, E. (2001). Labelling patterns and object naming. *Developmental science*, *4*(1), 109–118. https://doi.org/http://dx.doi.org/10.1111/1467-7687.00155

Lee, N. (2020). Learning (im)possible number syncretisms: Investigating innate featural representations. *Proceedings of the 50th Annual Meeting of the North East Linguistic Society*.

Leibovich-Raveh, T., Lewis, D. J., Kadhim, S. A.-R., & Ansari, D. (2018). A new method for calculating individual subitizing ranges. *Journal of Numerical Cognition*, *4*(2), 429–447.

Maldonado, M., & Culbertson, J. (2021). Person of interest: Experimental investigations into the learnability of person systems. *Linguistic Inquiry*. https://doi.org/10.1162/ling_a_00406

Noyer, R. R. (1992). *Features, positions and affixes in autonomous morphological structure* [Doctoral dissertation, Massachusetts Institute of Technology].

Saldana, C., Herce, B., & Bickel, B. (2022). More or less unnatural: Semantic similarity shapes the learnability and cross-linguistic distribution of unnatural syncretism in morphological paradigms. *Open Mind*, *6*, 183–210.

Saldana, C., Kanampiu, P., & Culbertson, J. (2025). Gender comes first: Experimental evidence for the representation of gender closer to the noun stem than number across linguistic populations. *Linguistic Inquiry*, 1–75.

Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological monographs: General and applied*, *75*(13), 1. https://doi.org/http://dx.doi.org/10.1037/h0093825

Silverstein, M. (1976). Hierarchy of features and ergativity. In R. M. Dixon (Ed.), *Grammatical categories in australian lan- guages* (pp. 112–171). Canberra: Australian Institutes of Aboriginal Stud- ies.

Skirgård, H., Haynie, H. J., Blasi, D. E., Hammarström, H., Collins, J., Latarche, J. J., Lesage, J., Weber, T., Witzlack-Makarevich, A., Passmore, S., et al. (2023). Grambank reveals the importance of genealogical constraints on linguistic diversity and highlights the impact of language loss.

Science Advances, *9*(16), eadg6175. https://doi.org/https://doi.org/10.1126/sciadv.adg6175

Trick, L. M., & Pylyshyn, Z. W. (1994). Why are small and large numbers enumerated differently? a limited-capacity preattentive stage in vision. *Psychological review*, *101*(1), 80.

Xu, F., & Tenenbaum, J. B. (2007). Word learning as bayesian inference. *Psychological review*, *114*(2), 245. https://doi.org/http://dx.doi.org/10.1037/0033-295X.114.2.245