Evidence of universal language structure from speakers whose language violates it

Alexander Martin,[1*] **David Adger,**[2] **Klaus Abels,**[3] **Patrick Kanampiu,**[4] **Jennifer Culbertson**[4]

[1]Center for Language and Cognition, University of Groningen

[2]Department of Linguistics, Queen Mary University of London

[3]Department of Linguistics, University College London

[4]School of Philosophy, Psychology and Language Sciences, The University of Edinburgh

[*]To whom correspondence should be addressed; E-mail: alexander.martin@rug.nl

Abstract

There is a longstanding debate in cognitive science surrounding the source of commonalities among languages of the world. Indeed, there are many potential explanations for such commonalities—accidents of history, common processes of language change, memory limitations, constraints on linguistic representations, etc. Recent research has used psycholinguistic experiments to provide empirical evidence linking common linguistic patterns to specific features of human cognition, but these experiments tend to use English speakers, who in many cases have direct experience with precisely the common patterns of interest. Here, we highlight the importance of testing populations whose languages go *against* cross-linguistic trends. We investigate whether monolingual speakers of Kîîtharaka, which has an unusual way of ordering words, mirror those of English speakers. We find that they do, supporting the hypothesis that universal cognitive representations play a role in shaping word order.

Statement of relevance

Claims of universality are commonly made in cognitive science, and they abound when it comes to language. For example, linguists appeal to universality to explain why certain word ordering patterns are found much more often than others across languages. Yet experimental evidence for universal representations or preferences tends to come exclusively from speakers whose languages *follow* cross-linguistic trends. These speakers are often from WEIRD (Western, Educated, Industrialized, Rich, Democratic) populations that are more accessible to researchers. But there is no guarantee that evidence from WEIRD populations will generalise to other populations. Indeed, there is good reason to suspect that speakers of languages that do not follow cross-linguistic trends will behave very differently from speakers of languages that do. Here, we compare two such populations, to test a hypothesized universal representation concerning word order. We find, perhaps surprisingly, that both populations behave similarly, providing strong support for this particular universal.

Evidence of universal language structure from speakers whose language violates it

Introduction

While human languages vary dramatically, they also seem to have certain properties in common. For example, in the domain of morphology, there is a robust tendency for languages to form complex words by adding suffixes to the ends of words—e.g., 'cat-s' or 'walk-ed'—rather than by adding prefixes to the beginnings—e.g., 'un-happy' or 'dis-trust'. In the domain of word order, languages tend to place adjectives (e.g., 'red') closer to the nouns they modify compared to demonstratives (e.g., 'that'). These kinds of commonalities are often called cross-linguistic generalizations, and they are extensively studied in the language sciences. What exactly explains them, however, is a source of ongoing debate (Bybee, 2008; Chomsky, 1995; Evans & Levinson, 2009; Prince & Smolensky, 1993/2004; Rooryck, Smith, Liptak, & Blakemore, 2010, a.o.). One possibility is that these commonalities reflect universal features, preferences, or biases of human cognition. However, the current statistical distribution of patterns across languages reflects many things, including which patterns emerged during the origins of human language (Gell-Mann & Ruhlen, 2011), accidents of history in the spread and death of languages since (Dunn, Greenhill, Levinson, & Gray, 2011), and processes of language change that are independent of cognition (Bybee, 2008). Therefore, whether a particular skewed distribution of linguistic patterns provides evidence for a cognitive universal is not trivial to establish (Chomsky, 2013; Culbertson, 2023; Evans & Levinson, 2009; Ladd, Roberts, & Dediu, 2014; Piantadosi & Gibson, 2014). Some evidence for a causal link between cross-linguistic generalizations and cognitive universals has come from experiments using artificial languages (Culbertson, 2023). Indeed, in a number of cases, participants more readily learn, generalise, or even alter the systems they learn to mirror common linguistic patterns (e.g., Culbertson & Adger, 2014; Fedzechkina, Jaeger, & Newport, 2012; Martin & White, 2021).

For example, the so-called suffixing bias—the cross-linguistic generalization that suffixing is more common than prefixing—has been argued to reflect a universal psychological principle: speakers attend more to word beginnings, and prefer to put lexical content there (e.g., the content word *cat*), and grammatical content (e.g., the plural marker *-s*), at the end (Hawkins & Cutler, 1988; Pycha, 2015). In line with this, evidence from artificial language experiments suggests that people treat novel words which differ in the presence of a suffix as

more similar to one another than novel words that differ in the presence of a prefix (Bruening, Brooks, Alfieri, Kempe, & Dabasinskiene, 2012; Hupp, Sloutsky, & Culicover, 2009). For example, given a novel word 'tate', participants judge 'tatebo' as more similar to it than 'botate'. Related studies suggest that that people may learn linguistic categories better when they are cued by a suffix as opposed to a prefix (St. Clair, Monaghan, & Ramscar, 2009).

In a similar vein, research studying a proposed universal bias in word ordering has also shown that the cross-linguistic trend for placing adjectives closer to nouns than demonstratives is also recapitulated in artificial language experiments. Participants taught nouns (like 'cup') and modifiers (like 'red' or 'that') will assume, without explicit evidence, that adjectival modifiers like 'red' should be ordered closer to nouns than demonstratives like 'that' (e.g., 'that red cup' or 'cup red that') (Culbertson & Adger, 2014; Martin, Abels, Ratitamkul, & Culbertson, 2019; Martin, Holtz, Abels, Adger, & Culbertson, 2020). This has been argued to support a potential cognitive universal favouring such orders.

Importantly however, these and most other such experiments target a very small subset of speaker populations, most often speakers of English and related languages. In addition to concerns about the lack of cultural variety more generally in experimental work (Henrich, Heine, & Norenzayan, 2010), the range of linguistic variety here is not sufficient to test the relevant hypotheses effectively (Blasi, Henrich, Adamou, Kemmerer, & Majid, 2022). In the case of the so-called suffixing preference, testing speakers of English cannot tell us much. English is a predominantly suffixing language itself, and speakers of English thus have extensive experience with a language which *adheres* to the relevant cross-linguistic generalization. It is therefore unsurprising to find, for example, that English speakers treat suffixed words as more similar to each other than prefixed words. Similarly, research on the order of modifiers in noun phrases has been conducted on speakers of English and Thai, which feature ordering patterns that conform to the relevant cross-linguistic generalization.

But in both cases, what is needed is evidence from populations whose experience diverges from the cross-linguistic trends in question—for example, speakers of a prefixing language, or speakers of a language with demonstratives closer to the noun than adjectives. By definition, languages which *violate* cross-linguistic generalisations are relatively rare, and sometimes extremely rare, and thus accessing participant populations who speak these languages can be very challenging. However, Majid (2023) highlights the importance of

theory-based sampling of this kind, arguing that claims of cognitive or psychological universality require experimentalists to test participants from cultures predicted to differ (or not) based on existing theories. Though large-scale comparative study may be the ideal, Majid (2023, p. 200) argues that when resources are limited (e.g., relevant populations cannot be, or are very difficult to access), "a critical test of universality can come from only two cultures–if those cultures are maximally distinct for the research question at hand". In the case of a hypothesised suffixing bias, comparing a pair of populations in this way has in fact revealed *distinct* preferences, not universality.

In a recent study, Martin and Culbertson (2020) compared English-speaking participants to speakers of Kîîtharaka, a Bantu language spoken in Kenya. In contrast to many of the world's languages, Kîîtharaka shows a strong tendency to form complex words using prefixes. The cognitive universal hypothesized above would nevertheless predict that these speakers will share their perceptions of complex words with English speakers. However the results of the study found that, unlike English speakers, Kîîtharaka speakers treated words which differed *in a prefix* as more similar than those which differ in a suffix. This suggests that speakers' perceptions are shaped by their experience with language: English speakers are used to processing words with suffixes; Kîîtharaka speakers are used to processing words with prefixes; thus across these populations, speakers' representations of similarity reflect their own experience. The results of this study suggest that when it comes to the cross-linguistic trend for suffixing, the best explanation is likely *not* a cognitive universal, but rather processes of language history and change (Bybee, Pagliuca, & Perkins, 1990; Enrique-Arias, 2002; Himmelmann, 2014).

In the remainder of this paper, we report a new study, using this same theoretically-motivated sampling method to study the proposed cognitive universal described above for the order of nouns and modifiers. We again compare speakers of a language which follows the cross-linguistic generalisation (English speakers) with speakers of a language which goes against it (Kîîtharaka speakers). In this case, we find evidence *supporting* the hypothesized universal.

A hypothesized universal in the domain of word order

Word order patterns are a rich domain in which to investigate commonalities among the world's languages and have drawn much attention in the linguistics and cognitive science literature. One extremely well-studied case is word order in complex noun phrases, like 'those two beautiful kittens', where a noun is combined with a series of modifiers. While there are 24 ways of ordering the four elements in this type of phrase (a demonstrative, a numeral, an adjective, and a noun), some orders are much more frequent than others (Cinque, 2005; Dryer, 2018; Greenberg, 1963, see Fig. 1A). An influential theory claims that noun phrase structure across languages involves a common hierarchical representation that reflects how human beings organise meaning and thought in the conceptual domain of entities (Rijkhoff, 2004). This hierarchical representation proposes that adjectives are most closely connected to nouns, while demonstratives are least closely connected (Abels & Neeleman, 2012; Adger, 2003; Alexiadou, Haegeman, & Stavrou, 2007; Cinque, 2005, see Culbertson, Schouwstra, and Kirby (2020) for discussion of the origins of the hierarchy itself). Some ways of linearly ordering elements in a noun phrase transparently reflect this hierarchy, and others do not. For example, the two most common orders, Noun Adjective Numeral Demonstrative (as in Thai), and Demonstrative Numeral Adjective Noun (as in English), reflect the hierarchy perfectly—adjectives are linearly closest to nouns, and demonstratives farthest away. But there are six additional orders that involve a transparent mapping from the hierarchy to linear order—i.e., they are *homomorphic* to the hierarchy (Fig. 1B)—which are also highly frequent. Together, over 80% of the world's languages have one of the eight homomorphic noun phrase word orders. A universal cognitive preference, or bias for homomorphism to the proposed underlying hierarchy is thus one hypothesised explanation for the skewed distribution.

In a series of experiments, Culbertson and Adger (2014) and Martin et al. (2020) found that English speakers who were taught miniature artificial languages were biased in favour of orders that were linearly different from English but which were nevertheless homomorphic. For example, they preferred homomorphic orders like Noun Adjective Demonstrative, over non-homomorphic orders like Noun Demonstrative Adjective, even though the latter has the same linear order of modifiers found in English (e.g., in 'these red cars', the demonstrative 'these' precedes the adjective 'red'). Martin et al. (2019) replicated this result for Thai speakers, showing that these speakers too ignored the linear order of modifiers in their native

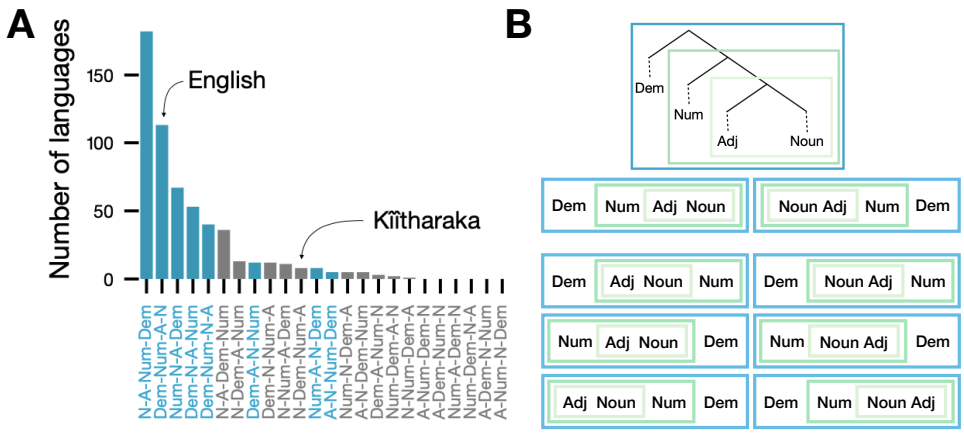

Figure 1. **A:** Distribution of noun phrase word orders across languages (blue patterns are homomorphic, data taken from Dryer, 2018). **B:** Hypothesised universal hierarchy. Nouns form a unit with adjectives, this unit combines with numerals, and this larger unit combines with demonstratives. Boxes immediately underneath the tree structure show the two linear orders that perfectly reflect the hierarchy (the most common), below that are the six other homomorphic orders.

language in favour of an order that maintains a homomorphic mapping to the hierarchy. These results have important implications for how language is represented in the mind. In particular, they illustrate that language is represented in a structured, hierarchical way, with elements recursively grouped together. Moreover, hierarchical representations impact how people generalize when they are learning a language: at least in this case, they are more important than linear order (Culbertson & Adger, 2014). In addition, the results are consistent with the cognitive universal proposed above: the representation that English and Thai speakers seem to use in these experiments could be a universal representation. However, English and Thai speakers' behaviour could also simply reflect prior experience with their own languages: because English and Thai are themselves homomorphic, they already provide structural evidence that adjectives should be closer to nouns and demonstratives farther away. Thus, evidence for universality must come from speakers of one of the rare languages whose noun phrase word order is *not homomorphic*. For these speakers, the hierarchy might look completely different: based on linear order, for example, demonstratives could be grouped with nouns more closely than adjectives.

The population we target here is Kîîtharaka speakers, the same population tested in

Martin and Culbertson (2020). As mentioned above, Kîîtharaka is a Bantu language spoken Kenya. In addition to being predominantly prefixing, Kîîtharaka is also unusual in having the demonstrative, rather than the adjective, closest to the noun: Noun Demonstrative Numeral Adjective, e.g., *tûbaka tûtû twîrî tûthongi*, *lit.* 'kittens those two beautiful'. If English speakers' bias for homomorphism is a result of their experience with their own language (which is itself homomorphic), then by extension, Kîîtharaka speakers should be expected to show the opposite preference—i.e., a bias for non-homomorphic orders. This is precisely what was found in previous experiments on suffixing: English speakers' behaviour (like their language) conformed to the cross-linguistic trend, but Kîîtharaka speakers showed the opposite preference (Martin & Culbertson, 2020). If Kîîtharaka speakers' preferences concerning order in the noun phrase align with English speakers', even though their own language is not homomorphic, that would provide the strongest type of evidence that this bias—and the hierarchical representation it reflects—are in fact universal.

Crucially, unlike Martin and Culbertson (2020), here we test *monolingual* Kîîtharaka speakers. These speakers live in relatively isolated areas of rural Tharaka-Nithi county (see Fig. 2A), and have no experience with other languages, like English, that could influence their use, or not, of a homomorphic order. We trained these monolingual Kîîtharaka speakers on an artificial language consisting of nouns, adjectives, and demonstratives (see Methods below). Participants learned this lexicon, and how to form simple phrases involving a single modifier and a noun. We taught participants that modifiers preceded the noun. This is the opposite of what is found in Kîîtharaka, where modifiers come after nouns. They heard phrases like *taka iti*, *lit.* 'red cup' (describing a neutrally-positioned red cup), or *himi iti*, 'this cup' (describing a grey-scale cup, spatially close to the speaker, see Fig. 2B).

Then, we showed participants new images whose descriptions required using both an adjective and a demonstrative, e.g., a red cup close to the speaker, or a black cup far from the speaker (see example test trial in Fig. 2B). Since participants were given no information about how the two types of modifiers should be ordered relative to one another, they simply had to take a guess. If there is a universal bias, then Kîîtharaka speakers, like English and Thai speakers, should guess a homomorphic order, with adjectives closest to the noun and demonstratives farthest away. If there is no universal bias, and instead the structural distance between nouns and modifiers in participants' native language determines their inferred order,

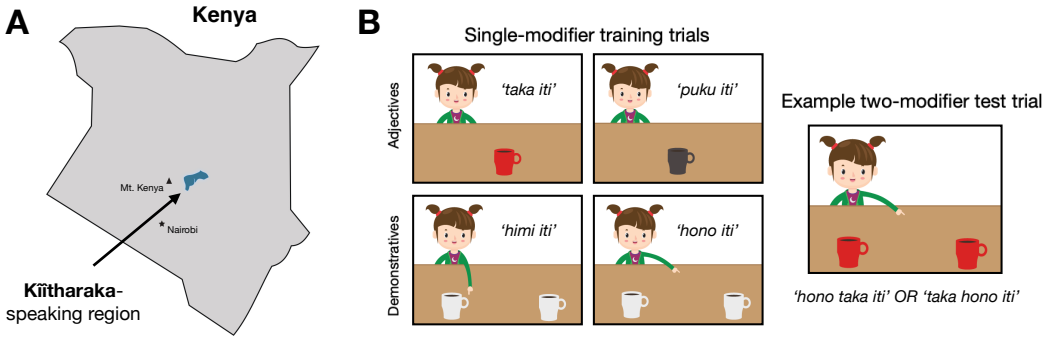

Figure 2. **A:** Kîîtharaka-speaking region in Kenya. **B:** Training stimuli (single modifier, either adjective or demonstrative). Example testing trial with both adjective and demonstrative, target meaning 'that red cup'.

then Kîîtharaka speakers should follow the structural closeness pattern in their own language and guess that the demonstrative comes closest to the noun.

There is a third possibility: that Kîîtharaka speakers will follow the *linear* order of modifiers in their language. This would lead them to produce Demonstrative Adjective Noun—superficially a homomorphic order, but potentially generated by simply following the surface modifier order of Kîîtharaka. We can confidently rule out this possibility through exactly the comparison we target—a homomorphic language like English, and a non-homomorphic language like Kîîtharaka. English speakers (and Thai speakers) have been consistently shown to ignore the linear order of modifiers in their native language in this task. In order to produce the homomorphic order they do, these speakers necessarily have to ignore this—they must invert the order of Demonstrative Adjective (Noun) to get (Noun) Adjective Demonstrative. Thus we have good reason to believe that what drives behaviour in this task is *not* linear order but structural order of modifiers (Culbertson & Adger, 2014; Martin et al., 2019, 2020). If the structural order English speakers used came from their native language, then we would expect Kîîtharaka speakers to produce a *non-homomorphic order*. If they produce a homomorphic order, then the most parsimonious explanation is that both populations are ignoring linear order and accessing a shared universal representation.

Methods

Anonymised coded data along with the full data cooking and analysis notebook and extended methods description can be found in the supporting information at the following

link: https://osf.io/xavb7/?view_only=8b364655f67740b2aab3e0a9af08a70e

Materials

The artificial language included three nouns (*eyey* 'feather', *uhu* 'ball', and *iti* 'cup'), two adjectives (*taka* 'red', and *puku* 'black'), and two demonstratives (*himi* 'this' and *hono* 'that'). Words were individually recorded by a phonetically-trained speaker. The visual stimuli consisted of cartoon images depicting different objects on a table in front of a cartoon girl (see Fig. 2B).

Procedure

The procedure followed Martin et al. (2020), but was adapted for this population by spreading the training and testing over two days. All participant recruitment and testing was done in Kîîtharaka by a local team member. On the first day, participants were trained on the nouns, as well as on combinations of a noun and a single modifier (adjective or demonstrative). Training consisted of passive exposure, picture matching (where two images appeared, a word or phrase in the language was played, and participants had to point to the corresponding image), and production (with feedback); see extended methodology in SI for details. On the second day, participants went through the training again, and then completed the critical testing block. Participants were told that they would have to describe images that they had not seen, using three words. They were shown an image depicting an object that was either red or black and was either in a proximal or distal position relative to a cartoon girl. There were 16 such trials. The noun was always "cup" (in order to ease the burden of lexical access). If participants had trouble remembering any lexical items, the experimenter could assist them with only one of them on a given trial.

Participants

91 participants were recruited from the Kîîtharaka-speaking region of Tharaka-Nithi County in Kenya between September 2019 and August 2022. In order to avoid any meaningful exposure to English, we recruited participants from rural areas who had little or no formal education. Our participants were thus older than those tested in typical (artificial language learning) experiments: median age = 48 years, maximum age = 79 years. We used strict inclusion criteria in order to analyse data only from participants who were verifiably

functionally monolingual. We excluded data from participants who self-reported any more than minimal English knowledge, or who were able to describe the two-modifier images using English words ($N = 25$). Data from participants who failed to produce at least 10 on-task responses (i.e., Demonstrative Adjective Noun or Adjective Demonstrative Noun) in the two-modifier phase ($N = 32$), or who were provided with more than one lexical item by the experimenter on any trial were also excluded ($N = 14$). We therefore analysed data from 20 monolingual participants.

Results

Focussing on the critical trials described above, we measured how often participants in our task produced a homomorphic order (1) and how often they produced a non-homomorphic order (0). Results are visualized in Fig. 3. We conducted our statistical analysis using logistic mixed-effects models implemented using the lme4 package (Bates, 2010) in R (R Core Team, 2020). We designed a full model with the binary dependent variable Homomorphic, along with by-participant random intercepts. We then used a likelihood ratio test to compare this model to a null model with no intercept term to test if, on average, participants chose homomorphic orders above the 50% chance level. We found a statistically significant difference between the full and null models ($\beta = 2.94$, $SE = 1.3$, $\chi^2(1) = 7.01$, $p < 0.001$), indicating an above-chance preference for homomorphic orders. We also compared the preferences of our Kîîtharaka-speaking participants with those of the English-speaking participants from Martin et al. (2020). We designed a full model with the binary dependent variable Homomorphic along with a deviation-coded factor Population and by-participant random intercepts, and compared this model to a simpler model excluding the Population factor. We found no significant difference between the two models ($\chi^2(1) < 1$) and thus no statistical evidence of a difference in preferences between the English- and Kîîtharaka-speaking samples. Given that there is no reason to assume Kîîtharaka speakers are different from English and Thai speakers in using structure rather than linear order to make inferences about a new language, this result supports the presence of a universal cognitive bias towards homomorphic orders in speakers of a language which goes against that bias.

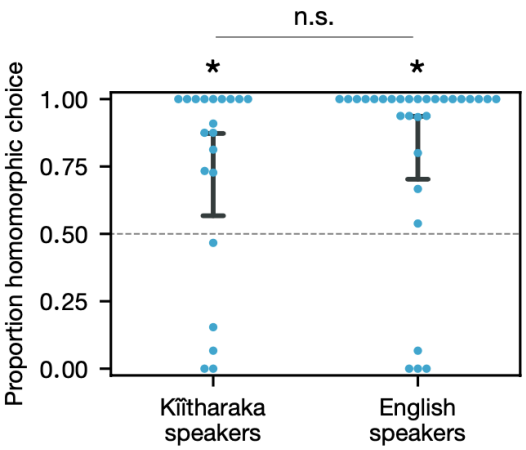

Figure 3. Experimental results showing consistent homomorphism preference across the two speaker populations (English results reproduced from Experiment 3 of Martin et al., 2020).

Discussion

Cross-linguistic generalisations have long provoked debates in linguistics and cognitive science: they represent intriguing commonalities across languages that might reflect cognitive universals, but alternative, non-cognitive explanations are also possible, and evidence adjudicating amongst the alternatives is not trivial to provide. Recent work has used methods from experimental psychology, like artificial language learning, to show that learners prefer linguistically common patterns. This alignment between learners' preferences and cross-linguistic generalisations supports the idea that the latter are driven by cognitive universals. However, like most psychological research, participant populations tend to come from a very narrow sample of the world's cultures and are overwhelmingly English-speaking. While English undoubtedly exhibits some unusual linguistic features (Blasi et al., 2022), in many cases, it conforms to cross-linguistic generalisations, raising the possibility that at least some previous results reflect English speakers' biases, not universal biases. Here, we compare data from English speakers tested in previous work to new data from monolingual Kîîtharaka speakers, in order to revisit a hypothesised explanation for a cross-linguistic generalisation about word order in complex noun phrases. These two populations differ in a crucial way: English conforms to the generalisation, and Kîîtharaka violates it. These two populations thus allow us to adjudicate between two different types of hypotheses for the cross-linguistic generalization in question: a cognitive universal, or cognition-external forces (like accidents of history, processes of change, etc.).

We found that despite differences in surface word order, the preferences of these two populations align. Both English and Kîîtharaka speakers prefer orders in which the adjective comes closest to the noun and the demonstrative farthest away. This is by far the most common type of pattern found across languages, but *not* in Kîîtharaka. That this population nevertheless prefers this kind of order when learning a new language is striking. It contrasts clearly with previous results showing that Kîîtharaka speakers do *not* show the same preference as English speakers in other domains of language (Martin & Culbertson, 2020). Our results therefore suggest that a cognitive universal drives the distribution of noun phrase word order patterns in the world's languages. Specifically, we found evidence for a universal bias reflecting a common underlying hierarchical structure – a representation that is shared across speakers, regardless of their native language. The origins of this shared representation remain a topic of ongoing investigation. Some researchers have argued that the hierarchical structure underlying nominal word order (and perhaps word order in general) is innate (Adger, 2003; Cinque, 2005). Other researchers have argued that it reflects conceptual knowledge about the world (e.g., that properties conveyed by adjectives are conceptually more closely associated with particular nouns than numerals or demonstratives are Culbertson et al., 2020), or perhaps relatedly, knowledge about which linguistic categories tend to be more informative about each other (Hahn, Degen, & Futrell, 2021). Regardless of how the hierarchy comes to be represented in speakers' minds, our findings suggest that the explanation must appeal to common cognitive mechanisms or experiences. Hierarchical structure has been argued to be one of the core features that makes human language special, and understanding the nature of these representations, whether they are shared, and how they are deployed is thus a fundamental question for the cognitive science of language.

More generally, our findings highlight the importance of evaluating explanations for common linguistic patterns with populations whose languages go *against* the trend. Alternative, non-cognitive explanations for universals—including accidents of history, random drift, common processes of language change—cannot be ruled out based on cross-linguistic data alone. But they also cannot be ruled out based on experimental evidence from English or a small sample of related languages; such experimental results cannot be assumed to replicate in all populations. Without evidence from diverse groups of learners whose experience differs in critical ways, it is impossible to make progress on fundamental questions about variation

and universality in our species.

References

Abels, K., & Neeleman, A. (2012). Linear asymmetries and the LCA. *Syntax*, *15*(1), 25–74.

Adger, D. (2003). *Core syntax*. Oxford: Oxford University Press.

Alexiadou, A., Haegeman, L., & Stavrou, M. (2007). *Noun Phrase in the generative perspective*. Berlin: Mouton de Gruyter.

Bates, D. (2010). lme4: Mixed-effects modeling with R. *http://lme4.r-forge.r-project.org/book*.

Blasi, D. E., Henrich, J., Adamou, E., Kemmerer, D., & Majid, A. (2022). Over-reliance on english hinders cognitive science. *Trends in cognitive sciences*.

Bruening, P., Brooks, P., Alfieri, L., Kempe, V., & Dabasinskiene, I. (2012). Children's tolerance of word-form variation. *Child Development Research*.

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

Bybee, J. L., Pagliuca, W., & Perkins, R. D. (1990). On the asymmetries in the affixation of grammatical material. In W. Croft & K. M. D. an Suzanne Kemmer (Eds.), *Studies in typology and diachrony: Papers presented to joseph h. greenberg on his 75th birthday* (pp. 1–42). John Benjamins Amsterdam.

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

Chomsky, N. (2013). Problems of projection. *Lingua*, *130*, 33 - 49. Retrieved from `http://www.sciencedirect.com/science/article/pii/S002438411200277X` doi: http://dx.doi.org/10.1016/j.lingua.2012.12.003

Cinque, G. (2005). Deriving Greenberg's Universal 20 and its exceptions. *Linguistic Inquiry*, *36*(3), 315-332.

Culbertson, J. (2023). Artificial language learning. In J. Sprouse (Ed.), *Oxford handbook of experimental syntax.* New York, NY: Oxford University Press.

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

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).

Dryer, M. (2018). On the order of demonstrative, numeral, adjective and noun. *Language*, *94*(4), 798–833.

Dunn, M., Greenhill, S., Levinson, S., & Gray, R. (2011). Evolved structure of language shows lineage-specific trends in word-order universals. *Nature*, *473*(7345), 79–82.

Enrique-Arias, A. (2002). Accounting for the position of verbal agreement morphology with psycholinguistic and diachronic explanatory factors. *Studies in Language. International Journal sponsored by the Foundation "Foundations of Language"*, *26*(1), 1–31.

Evans, N., & Levinson, S. C. (2009). The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences*, *32*(05), 429-448.

Fedzechkina, M., Jaeger, T. F., & Newport, E. L. (2012). Language learners restructure their input to facilitate efficient communication. *Proceedings of the National Academy of Sciences*, *109*(44), 17897–17902.

Gell-Mann, M., & Ruhlen, M. (2011). The origin and evolution of word order. *Proceedings of the National Academy of Sciences*, *108*(42), 17290-17295.

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*.

Hawkins, J. A., & Cutler, A. (1988). Psycholinguistic factors in morphological asymmetry. In J. A. Hawkins (Ed.), *Explaining language universals* (pp. 280–317). Oxford: Blackwell.

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, *33*(2-3), 61–83.

Himmelmann, N. P. (2014). Asymmetries in the prosodic phrasing of function words: Another look at the suffixing preference. *Language*, *90*(4), 927–960.

Hupp, J. M., Sloutsky, V. M., & Culicover, P. W. (2009). Evidence for a domain-general mechanism underlying the suffixation preference in language. *Language and Cognitive Processes*, *24*(6), 876-909.

Ladd, D. R., Roberts, S. G., & Dediu, D. (2014). Correlational studies in typological and historical linguistics. *Annual Review of Linguistics*, *1*.

Majid, A. (2023). Establishing psychological universals. *Nature Reviews Psychology*, *2*(4),

199–200.

Martin, A., Abels, K., Ratitamkul, T., & Culbertson, J. (2019). Cross-linguistic evidence for cognitive universals in the noun phrase. *Linguistics Vanguard*, *5*(1).

Martin, A., & Culbertson, J. (2020). Revisiting the suffixing preference: Native language affixation patterns influence perception of sequences. *Psychological Science*, *31*(9), 1107–1116.

Martin, A., Holtz, A., Abels, K., Adger, D., & Culbertson, J. (2020). Experimental evidence for the influence of structure and meaning on linear order in the noun phrase. *Glossa*, *5*(1), 1-21.

Martin, A., & White, J. (2021). Vowel harmony and disharmony are not equivalent in learning. *Linguistic Inquiry*, *52*(1), 227–239.

Piantadosi, S. T., & Gibson, E. (2014). Quantitative standards for absolute linguistic universals. *Cognitive Science*, *38*, 736–756.

Prince, A., & Smolensky, P. (1993/2004). *Optimality theory: Constraint interaction in generative grammar*. New York, NY.

Pycha, A. (2015). Listeners perceive prefixes differently: Evidence from a noise-rating task. *Word Structure*, *8*(1), 53–83.

R Core Team. (2020). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria.

Rijkhoff, J. (2004). *The noun phrase*. Oxford: Oxford University Press.

Rooryck, J., Smith, N., Liptak, A., & Blakemore, D. (2010). Special issue on Evans & Levinson's "The myth of language universals". *Lingua*, *120*, 2651–2758.

St. Clair, M. C., Monaghan, P., & Ramscar, M. (2009). Relationships between language structure and language learning: The suffixing preference and grammatical categorization. *Cognitive Science*, *33*, 1317–1329.