Syntactic harmony arises from a domain-general learning bias

Syntactic harmony arises from a domain-general learning bias

Jennifer Culbertson (jennifer.culbertson@ed.ac.uk)
Centre for Language Evolution, 3 Charles Street
Edinburgh, UK EH8 9AD

Simon Kirby (simon.kirby@ed.ac.uk)
Centre for Language Evolution, 3 Charles Street
Edinburgh, UK EH8 9AD

Abstract

Syntactic harmony occurs when heads and dependents align within and across different types of phrases in a language. Harmony is a well-known (statistical) typological universal: in most languages, many if not all heads and dependents are consistently ordered (i.e., either head-dependent, or dependent-head). Despite decades of work, from every conceivable theoretical perspective, the origins of syntactic harmony remain opaque. However, recent work using artificial language learning has suggested that harmonic patterns are *easier to learn* than their non-harmonic counter-parts. Thus at least part of the explanation for this tendency may be linked to learning. Here, we explore whether the mechanism behind the learning bias for syntactic harmony is fundamentally domain-general by instantiating harmony in non-linguistic stimuli. Our findings support the claim that the origins of syntactic harmony lie in a domain-general bias for simplicity acting on linearized, language-specific categories.

Keywords: language universals; syntax; cognition; learning biases; artificial grammar learning

Introduction

Background

Word order harmony occurs when the order of syntactic heads and their dependents follows a consistent pattern within a given language. For example, SOV languages are head-final in the verb phrase. In other words, the syntactic head, the verb, follows its dependent, the object. These languages also tend to have postpositions, where the head, the adposition, follows its dependent, the noun phrase. Similarly, languages with pre-nominal adjectives, tend to have all other nominal modifiers occurring before the noun as well (e.g., numeral words, demonstratives). Harmony is one of the oldest, and most well-studied so-called 'typological universals', and was first documented at length by Greenberg (1963). Since then, many explanations for harmony have been proposed. Indeed, it has played a central role in one of the most important debates in linguistics, namely whether language is shaped in meaningful ways by features of the human cognitive and/or linguistic system.

While most researchers would likely concede the importance of general cognition in language, the existence of linguistic-specific factors that constrain language is much more contentious (e.g., see Evans & Levinson, 2009). In a recent paper, Culbertson and Kirby (2016) argue that there are (at least) two relevant ways in which linguistic-specific factors can influence some feature of language typology–like word order harmony. The most obvious way is if that feature evolves by natural selection under a pressure for the linguistic function it serves. The second, and perhaps less obvious, way is if that feature is cognition-general but *interacts* with the linguistic system in a unique way. Importantly, the implications for evolution are radically different for these two paths. Computational models of language evolution suggest that the first situation is high unlikely (e.g., see Chater, Reali, & Christiansen, 2009; Smith & Kirby, 2008). However, the second situation is much more plausible, and indeed Culbertson and Kirby (2016) argue that harmony represents such a case.

Specifically, Culbertson and Kirby (2016) argue that harmony is one particular instantiation of a cognition-general bias for simplicity in learning (Chater & Vitányi, 2003). A simplicity bias can be expressed as a preference for inferring explanations (e.g., grammars) that are simpler, or put another way, can be described more concisely (in information-theoretic terms). A simplicity bias *interacting with* a linguistic-specific categorization of grammatical elements into heads and dependents, can explain a preference for harmony. A language with a single, general rule governing the order of these categories across phrases is simpler than one with multiple, specific ordering rules applied to different types of phrases. This kind of explanation for harmony is related to early proposals like the Head-Direction Parameter (i.e., a high-level rule, set to head-initial or head-final in a given language Travis, 1984; Chomsky, 1988; Baker, 2001). However, these proposals posited innate linguistic-specific constraints, generating universally possible languages and ruling out impossible ones. By contrast, the simplicity bias is critically *not* specific to the linguistic system, and does not itself make predictions about *possible* grammars, but instead *likely* grammars.

While this view might be intuitively plausible, there are a number of well-known alternative explanations for harmony. The first is an alternative explanation that appeals to sentence processing. In particular, a number of researchers have argued that at least some harmonic orders minimize the dependency lengths of utterances consisting of multiple heads and dependents (e.g., Hawkins, 2004; Temperley & Gildea, 2018; Futrell, Mahowald, & Gibson, 2015; Hahn, Jurafsky, & Futrell, 2020). For example, a sentence like 'Ally sent the report to the cabinet secretary' involves harmonic VP and PP order, but also a shorter distance between the verb

In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.

'sent' and the preposition 'to' indicating a PP dependent compared to a hypothetical non-harmonic alternative like 'Ally sent the report the cabinet secretary to'. Like simplicity, dependency-length minimization is likely a domain-general cognitive pressure that applies to language (e.g., see Temperley & Gildea, 2018). The second alternative explanation originates with Givón (1975, 1979), who observed that harmonic heads were in some cases diachronically related. For example, he argued that the historical source for many adpositions is verbs[1], thus we expect adposition order to correlate with verb order. If harmony among particular heads were widely amenable to this sort of explanation, then the typological trend for harmony could thus reflect largely non-cognitive processes. In other words it could result not from individual-level cognitive biases, but purely from common historical processes of grammaticalization (see also Heine & Kuteva, 2007; Aristar, 1991; Whitman, 2008; Kaufman, 2009; Collins, 2019). A final alternative explanation is that harmony is a historical accident, with the typological trend for harmony resulting from the fact that a number of large language families happen to be harmony. For example, Dunn, Greenhill, Levinson, and Gray (2011) use Bayesian phylogenetic methods to evaluate whether historical changes to head-dependent order in one phrase correlate with changes in another phrase. If this happens consistently across language families then this implies a general pressure for harmony. Although they find evidence of correlated pairs of phrases within some of the language families, the results differ across families. Even correlations which appear to be very strong on the basis of Dryer (1992) are thus interpreted as 'lineage-specific': for example, the correlation between verb-object and adposition-noun order was found to be strong in the Indo-European and Austronesian families, but not in Bantu or Uto-Aztecan.

Harmony in behavioral experiments

In summary, there are a number of potential explanations for harmony that differ in the degree to which they posit linguistic-specific rules and representations (as in the Head-Direction Parameter and related proposals Baker, 2001), cognition-general mechanisms acting on linguistic representations (as in Culbertson & Kirby, 2016; Temperley & Gildea, 2018), or neither (as in Givón, 1975; Dunn et al., 2011). These explanations have generated robust debate for decades, without a clear answer, suggesting the need for additional sources of evidence beyond typological data.

Recent work has used artificial language learning experiments to more directly test the hypothesized link between cognition and harmony (see Culbertson, to appear, for a review). For example, Culbertson, Smolensky, and Legendre (2012) taught adult English-speaking learners a miniature artificial language featuring noun phrases consisting of either a noun and an adjective or a noun and a numeral. They found that learners were more successful at learning the language if it generally placed both modifiers on the same side of the noun. In other words, harmonic patterns were easier to learn (see also Culbertson & Newport, 2015, 2017; Culbertson, Franck, Braquet, Barrera Navarro, & Arnon, 2020, for evidence from children and speakers of other languages). Critically, this is in line with a simplicity account of harmony, or with something like a Head-Direction Parameter, but not with a non-cognitive accounts of harmony, nor with a dependency-length minimization account: in these experiments only a single phrase is processed (or produced) at a time, and therefore all input and output languages have the same (minimal) dependency lengths (for evidence in favor of dependency-length minimization in artificial language learning see Fedzechkina, Chu, & Florian Jaeger, 2018) These results suggest that a cognitive bias may play a causal role in explaining word order harmony. However, it remains unclear whether the bias is linguistic in nature, or reflects a cognition-general bias for simplicity.

Exploring the domain-general nature of harmony

Culbertson and Kirby (2016) propose that linguistic harmony reflects a domain-general bias for simplicity *in combination with* a linguistic-specific notion of similarity among elements. This notion of similarity is what determines which elements in the system should align with which others. In the case of language, this notion in question is the distinction between grammatical categories of words, i.e., different types of heads and dependents. While exactly how grammatical categories are defined and what determines whether a given category is a head or a dependent is necessarily straightforward, a specific theory of these categories will make clear predictions about which elements should align with which. In other words, the bias for simplicity will favour consistent ordering of those elements which are defined as similar in the relevant way. In this sense, harmony is the result of both cognition-general and linguistic-specific factors. Here we test this proposal by generating *non-linguistic* stimuli which also feature similarities among elements that are specific to the domain in question. As we discuss further below, the asymmetry between heads and dependents is less important for our purposes than the notion of similarities between distinct categories. Therefore, while will call these heads and dependents, to illustrate the parallel with syntactic harmony, these are essentially a linguistic notions. Below, we explore whether systems with consistent harmonic alignment of similar elements are easier to learn. If learners find such systems easier to learn, this supports the view of harmony driven by a general bias for simplicity, rather than a view which posits harmony as the result of a constraint (e.g., parameters) evolved by natural selection under a pressure for linguistic harmony.

Experiment 1

In Experiment 1 we take one step away from syntactic harmony in natural language by using sequences of meaningless letter strings rather than combinations of meaningful words.

[1] Specifically, serial verb constructions, where two verbs are essentially concatenated, and one is eventually re-interpreted as an adposition.

The design is an ease-of-learning experiment, where participants are taught either a harmonic or non-harmonic ordering and tested on how well they learn to identify correct sequences.

Methods

Stimuli Stimuli consisted of two categories: 'head' elements and 'dependent' elements. In natural language the elements that fall into these categories–i.e., grammatical categories like noun or adposition—are presumably learned based on similarities in structure and/or meaning. As mentioned above, what exactly determines what is a head and what is a dependent is determined by linguistic-theory-internal notions. Since we are not interested here in how the categories themselves are learned, our stimuli are designed to provide salient evidence of distinct categories, and cues to which go together. In particular, there are similarities among the elements in each category (akin to similarities that nouns sharing with other nouns, or verbs share with other verbs), and among the heads and dependents that go together (akin to something like feature-based agreement). The stimuli are shown in Table 1. Head categories were longer than dependent categories, and heads were presented in isolation at the beginning of the experiment and described to participants as having shorter sequences which attach to them (see Procedure below). In all other respects, what we are calling heads and dependents were mutually dependent on each other. Both head and dependent categories are made up of two types, with 4 tokens per type.[2] Heads types are distinguished by a contrast in voicing and manner of articulation of the relevant letters in the string. One type of head is comprised of CVCVC strings where C's are {n,g,ng} and V's are {a,e}; the other type of head has C's {sh,k,th} and V's are {u,o}. Tokens of each head type are distinguished by a different arrangement of the relevant set of consonants and vowels. Dependents are all comprised of CVC strings. Each head type can occur with two dependent types. Dependents paired with each type of head have C's that match it in terms of voicing, and V's chosen from the same set in order to help learner identify the relevant head-dependent combinations (akin to different phrase types in natural language): {b,v} with {a} or {e} for one type; {p,f} with {o} or {u} for the other.[3] Strings were constructed by appending dependents before or afters heads with a dash connecting them (see Table 2).

Participants were randomly assigned to one of three conditions which differed in the ordering of heads and dependents: harmonic, non-harmonic across heads, and non-harmonic within heads, all illustrated in Table 2. In the harmonic condition all string were consistently ordered. Participants in the harmonic condition were randomly assigned to H(ead)-Dep(endent) or Dep-Head order. In the non-harmonic across heads condition, all dependents of a particular head type were consistently ordered, but that order differed across head types. For example, in natural language, verbs might always precede their dependents, but adpositions might always follow their dependents. Participants in the non-harmonic across heads condition were randomly assigned to H1-Dep1, Dep2-H2 or the reverse. In the non-harmonic within heads condition, the two dependent types of a given head were inconsistently ordered. For example, Dep1a-H1, H1-Dep1b, Dep2a-H2, H2-Dep2b. For example, in natural language, adjectives might come before the head noun, but numerals after, and for some other head type, there might be dependents that proceed and others that follow. There are four possible combinations of orders in this condition, and participants are randomly assigned to one of them. Our prediction is that participants in the harmonic condition will learn more accurately than participants in either non-harmonic condition. We further predict that the non-harmonic across heads pattern will be easier to learn than non-harmonic within heads pattern.

Table 1: Experiment 1 head and dependent stimuli.

Heads		Dependents
H1	{nageng, negang, genang, ganeng}	Dep1a: {bav, baz, dav, daz} Dep1b: {veb, ved, zeb, zed}
H2	{shukoth,shokuth, koshuth, kushoth}	Dep2a: {puf, pus, tuf, tus} Dep2b: {fop, fot, sop, sot}

Table 2: Experiment 1 conditions with example sequences from the first listed subcondition.

Condition	Sub-conditions	Example seq.
Harmonic	Head-Dep or Dep-Head	nageng-bav ganeng-veb shokuth-tuf koshuth-fop
Non-harmonic across heads	Dep-H1, H2-Dep or H1-Dep, Dep-H2	bav-nageng veb-ganeng shokuth-tuf koshuth-fop
Non-harmonic within heads	e.g., Dep1a-H1, H1-Dep1b Dep2a-H2, H2-Dep2b	bav-nageng ganeng-veb tuf-shokuth koshuth-fop

Procedure The experiment was presented in a web browser using jspysch (De Leeuw, 2015). Participants were informed that they would be learning to recognize two new types of letter sequences. Each new type of sequences has a shorter sequence associated with it, and their task was to learn to recognize how the short sequences attach to the long ones.

[2] In natural languages there may be asymmetries between the number of types and tokens for various categories of heads and dependents, but we have no particular reason to expect that this plays a role in harmony.

[3] There is nothing special about these choices of how to differentiate the categories, or how to indicate the heads and dependents that go together; these were just features that we thought would be highly salient for learners.

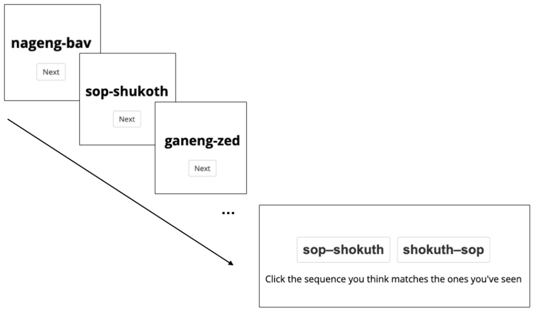

Figure 1: Example training and testing trials in Experiment 1.

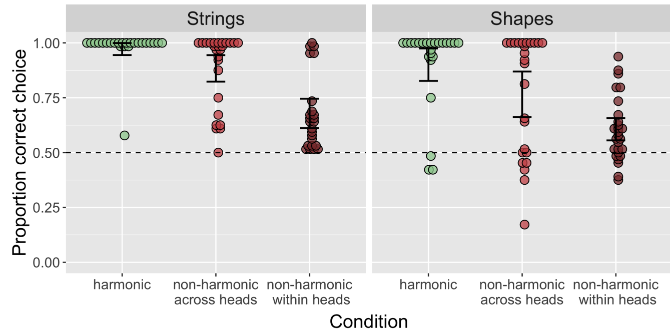

Figure 2: Proportion correct choice for each condition in Experiment 1 (Strings) and Experiment 2 (Shapes). In both cases, accuracy is highest in the harmonic condition and lowest in the non-harmonic within heads condition.

They were then shown examples of the two types of heads in isolation on the screen, with a border around the examples of each type. After this, they were trained on letter sequences one at a time. The sequence appeared, and after 1500ms a button labelled with 'Next' appeared and participants clicked to advance to the next trial. Training consisted of 64 trials (one repetition of each possible combination of head with dependent in the language). After training, participants were tested on what they had learned. In each testing trial, two sequences appeared on the screen. Participants were instructed to choose which was a possible configuration of the sequences they had learned about. Testing consisted of 64 trials (one repetition of each correct combination of head with dependent in the language, paired with a sequence which reversed the order). Example training and testing trials are shown in Figure 1.

Participants Participants were 74 English-speakers who self-identified as monolinguals on the Prolific Academic platform (24 harmonic, 26 non-harmonic across heads, 24 non-harmonic within heads).

Results

Recall that based on Culbertson and Kirby (2016) we predicted that a harmony bias should still be present when instantiated in sequences of meaningless letter strings. More specifically, we predicted that participants would have the highest accuracy in the harmonic condition, and the lowest accuracy in the non-harmonic within heads condition. Figure 2 suggests that this prediction was borne out. The data were analysed using mixed-effect logistic regression in order to assess the effect of condition on accuracy. We compared two models, one including condition as a predictor and the other include only an intercept term. Both models including a by-participant random intercept. A likelihood ratio test indicated that including condition significantly improved the model ($chi^2 = 51.07, p < 0.001$). We conducted a further model comparing performance in the harmonic to the two non-harmonic conditions. In both cases, accuracy was higher in the harmonic condition (vs. non-harmonic across heads $\beta = -3.18 \pm 0.81, p < 0.001$; vs. non-harmonic within heads $\beta = -5.70 \pm 0.83, p < 0.001$). Finally, we compared the two non-harmonic conditions, confirming that accuracy was higher in the non-harmonic across heads condition than the non-harmonic within heads condition ($\beta = -2.38 \pm 0.55, p < 0.001$).

Discussion

In Experiment 1, we tested the hypothesis that linguistic harmony is driven by a cognition-general simplicity bias combined with a linguistic-specific notion of heads and dependents which determines the elements-to-be-aligned. This hypothesis predicts that a preference for harmony should be found in a tasks that do not involve syntactic phrases, heads or dependents; in particular where the stimuli *and* and the relevant categories are not syntactic in nature. Here, the stimuli were meaningless letter strings, with the 'head' and 'dependent' categories distinguished based on letters and sounds. We found that, as predicted, participants were better at learning the system when trained on sequences in which those heads and dependents were harmonic. They were less successful when sequences involves inconsistent ordering across different types of heads, and they were even less successful when there was inconsistent order of the different dependents of a single type of head. In Experiment 2 we move further away from language, instantiating harmonic and non-harmonic patterns in shape stimuli.

Experiment 2

Experiment 2 was identical to Experiment 1 in terms of the structure of the training, and the experimental conditions participants were assigned to. The only difference was in how head and dependent elements were constructed to instantiate harmonic and non-harmonic patterns.

Methods

Stimuli Stimuli consisted of two categories of shapes: 'head' elements and 'dependent' shapes. The categories are distinguished based on visual characteristics including size and structure as shown in Table 3. Heads were larger than

Table 3: Head and dependent stimuli for Experiment 2.

Heads		Dependents
H1		Dep1a: 
		Dep1b: 
H2		Dep2a: 
		Dep2b: 

dependents, and (as for Experiment 1) were first presented in isolation and were described to participants as having smaller shapes attached to them. As in Experiment 1, both head and dependent categories were made up of two types, with 4 tokens per type. Heads types were distinguished by their shape. One type of head consisted of circular shapes with rounded interior structure, the other type consisted on octagonal shapes with angular interior structure. Tokens of each head type were distinguished by their distinct interior structure. Dependents were all comprised of two small shapes put together. Each head type could occur with two dependent types. Dependents paired with each type of head had round/angular shapes which matched the head in order to help learners identify the relevant head-dependent combinations: oval with skinny rounded flower and rounded rectangle with fat flower for one type; diamond with skinny star and rectangle with fat star for the other. Sequences were constructed by appending dependents before or afters heads with a smooth or jagged line connecting them (see Figure 3). Conditions were the same as in Experiment 1.

Procedure The procedure was identical to Experiment, except that participants were instructed that they would be learning about two new types of shapes which had smaller shapes attached to them. Example training and testing trials are shown in Figure 3.

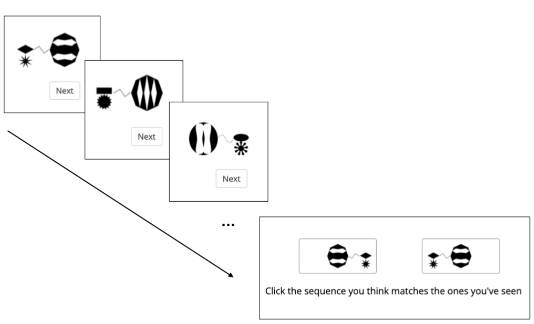

Figure 3: Example training and testing trials in Experiment 2.

Participants Participants were 76 English-speakers who self-identified as monolinguals on the Prolific Academic platform (24 harmonic, 25 non-harmonic across heads, 27 non-harmonic within heads).

Results

Recall that based on Culbertson and Kirby (2016) we predicted that a harmony bias should still be present when instantiated in non-linguistic stimuli, here sequences (or visual arrangements) of shapes. More specifically, we predicted that participants would have the highest accuracy in the harmonic condition, and the lowest accuracy in the non-harmonic within heads condition. As for Experiment 1, Figure 2 suggests that this prediction was borne out for these stimuli as well. The data were analysed using mixed-effect logistic regression in order to assess the effect of condition on accuracy. We compared two models, one including condition as a predictor and the other include only an intercept term. Both models including a by-participant random intercept. A likelihood ratio test indicated that including condition significantly improved the model ($chi^2 = 33.50, p < 0.001$). We conducted a further model comparing performance in the harmonic to the two non-harmonic conditions. In both cases, accuracy was higher in the harmonic condition (vs. non-harmonic across heads $\beta = -1.88 \pm 0.72, p < 0.001$; vs. non-harmonic within heads $\beta = -4.02 \pm 0.71, p < 0.001$). Finally, we compared the two non-harmonic conditions, confirming that accuracy was higher in the non-harmonic across heads condition than the non-harmonic within heads condition ($\beta = -1.92 \pm 0.51, p < 0.001$).

Discussion

In Experiment 2, we tested the hypothesis that linguistic harmony is driven by a cognition-general simplicity bias combined with a linguistic-specific notion of heads and dependents which determines the elements-to-be-aligned. This hypothesis predicts that a preference for harmony should be found in a tasks that involve completely non-linguistic categories as stand-ins for heads and dependents. Here, the stimuli were shapes, with the 'head' and 'dependent' categories distinguished based on size, shape, and structure. We found that, as predicted, participants were better at learning the system when trained on sequences in which those heads and dependents were harmonic. They were less successful when sequences involves inconsistent ordering across different types of heads, and they were even less successful when there was inconsistent order of the different dependents of a single type of head.

General Discussion

Word order harmony involves consistent alignment of heads and dependents within and across different phrases—for example, alignment of adpositionals and their noun dependents with verbs and their objects. Since the original observations of (Greenberg, 1963), linguists have sought to understand why languages tend to use harmonic orders. Traditionally, there have been several prominent competing hypotheses in the literature including a constraint, specific to syntactic system, which determines head order and includes a mechanism by which harmonic orders are preferred (as in Travis, 1984; Baker, 2001), a sentence processing mechanism along the lines of dependency-length minimization (as in Hawkins, 2004), and a cognition-external explanation based on patterns of language change (as in Givón, 1975; Aristar, 1991). A further possibility is set out in Culbertson and Kirby (2016), where harmony is hypothesized to result from a cognition-general bias for simplicity acting on linearized, language-specific categories. These alternative explanations have very different implications for language evolution. At least one implies a strong, domain-specific constraint, others imply no role for human cognition at all. A number of strands of research suggest that the idea of strong constraints which are specific to the linguistic system, and have evolved for a particular linguistic purpose–e.g., to enforce harmony–are very unlikely (Thompson, Kirby, & Smith, 2016). However, weak (i.e., defeasible) cognition-general biases like simplicity have been argued to be much more likely. Moreover, these bias can interact with unique features of our linguistic system, like syntactic categories, phrases, heads and dependents. However, direct behavioral evidence for such constraints is needed.

Here we have presented two experiments, inspired by previous research using artificial language learning methods showing that harmony among linguistic categories is preferred by learners. These findings suggest that there is indeed a link between individual-level biases of learners and a population-level tendency for harmony. We tested the prediction of Culbertson and Kirby (2016), that when the elements to be linearized are meaningless (sequences of letter strings) and/or non-linguistic (sequences of shapes), a simplicity bias should nevertheless lead to a preference for harmonic orders. This prediction was clearly borne out in both experiments.

It is worth noting again here that our design was based on the idea that the key aspect of harmony which links to simplicity is the linearization of similar categories of elements across sequences. In syntax, these are heads and dependents (defined in some way by a particular linguistic theory). Here we borrowed this terminology, and created distinct categories of elements (identifiable based on orthographic/phonological, or visual features) which could occur together. Learners were tasked with acquiring how these elements were linearized. In the harmonic condition, learners could in principle learn the correct linearization of elements without learning anything other than the distinction between the head and dependent categories–i.e., they could simply learn that dependents come first or last. This is part of what makes this type of pattern simple. In the non-harmonic across heads condition, by contrast, participants must make use of information about classes of dependents, or classes of heads, or both in order to take advantage of within-head consistency. In the non-harmonic within heads condition, yet more information about the heads and dependents must be learned in order to learn the linearization pattern. Future work could explore more deeply what exactly participants have learned about the dependencies in these very simplified systems, and the connection between harmony and category/dependency learning. However, we believe these experiments provide the first step toward showing that a weak cognition-general bias for simpler representations, active in individual learners, and amplified over time via cultural transmission, can explain syntactic harmony.

Of course, cognition-external factors like genetic relationship among languages, common grammaticalization pathways, etc. have also undoubtedly shaped language. The role of simplicity in driving harmony does not preclude the role of likely types of lexical changes (i.e., verbs being re-interpreted as adpositions). However, it does beg the question of whether the commonality of certain types of lexical changes on the one hand, and the strength of harmony between a given pair of categories on the other, are driven by the same underlying cause. For example, verbs and adpositions tend to be related by a common diachronic root *and* the tendency for harmony between them may be particular strong across languages (e.g. see, Dryer, 1992). One possibility is that syntactic harmony results not just from the linguistic-specific notions of head and dependent, but from more fine-grained representations of similarity among different types of grammatical categories. Categories of heads that are more similar to one another might be more likely to change into each other, and to harmonize. Whether this is the case remains to be determined through additional typological, theoretical, and experimental work.

Similarly, there is good evidence that languages are shaped by a preference for minimizing dependency-lengths during sentence processing (e.g, Temperley & Gildea, 2018; Futrell et al., 2015). However, the findings reported here, along with results from previous artificial language learning studies (e.g., Culbertson et al., 2012; Culbertson & Newport, 2015; Culbertson et al., 2020) suggest that harmony may be driven by a higher-level preference for representational simplicity which holds even when alignment only holds across utterances. How and whether these two forced interact–e.g., to strengthen the preference for harmony between types of phrases that frequently co-occur in a single utterance–is also an area for further inquiry.

Acknowledgments

This project was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 757643) and the Economic and Social Research Council (grant number ES/R011869/1).

References

Aristar, A. R. (1991). On diachronic sources and synchronic pattern: An investigation into the origin of linguistic universals. *Language*, *67*(1), 1-33.

Baker, M. (2001). *The atoms of language: The mind's hidden rules of grammar*. New York, NY: Basic Books.

Chater, N., Reali, F., & Christiansen, M. H. (2009). Restrictions on biological adaptation in language evolution. *Proceedings of the National Academy of Sciences*, *106*(4), 1015–1020.

Chater, N., & Vitányi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences*, *7*(1), 19–22.

Chomsky, N. (1988). *Language and problems of knowledge: The Managua lectures*. Cambridge, MA: MIT Press.

Collins, J. (2019). Some language universals are historical accidents. *Explanation in typology: Diachronic sources, functional motivations and the nature of the evidence*, *3*, 47.

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

Culbertson, J., Franck, J., Braquet, G., Barrera Navarro, M., & Arnon, I. (2020). A learning bias for word order harmony: evidence from speakers of non-harmonic languages. *Cognition*, *204*.

Culbertson, J., & Kirby, S. (2016). Simplicity and specificity in language: Domain general biases have domain specific effects. *Frontiers in Psychology*, *6*(1964).

Culbertson, J., & Newport, E. L. (2015). Harmonic biases in child learners: In support of language universals. *Cognition*, *139*, 71-82.

Culbertson, J., & Newport, E. L. (2017). Innovation of word order harmony across development. *Open Mind: Discoveries in Cognitive Science*, *1*(2), 91–100.

Culbertson, J., Smolensky, P., & Legendre, G. (2012). Learning biases predict a word order universal. *Cognition*, *122*, 306-329.

De Leeuw, J. R. (2015). jspsych: A javascript library for creating behavioral experiments in a web browser. *Behavior research methods*, *47*(1), 1–12.

Dryer, M. (1992). The Greenbergian word order correlations. *Language*, *68*(1), 81-183.

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.

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., Chu, B., & Florian Jaeger, T. (2018). Human information processing shapes language change. *Psychological science*, *29*(1), 72–82.

Futrell, R., Mahowald, K., & Gibson, E. (2015). Large-scale evidence of dependency length minimization in 37 languages. *Proceedings of the National Academy of Sciences*, *112*(33), 10336–10341.

Givón, T. (1975). Serial verbs and syntactic change: Niger-congo. In C. Li (Ed.), *Word order and word order change* (p. 47-112). Austin, TX: University of Texas Press.

Givón, T. (1979). *On understanding grammar*. New York: Academic Press.

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., Jurafsky, D., & Futrell, R. (2020). Universals of word order reflect optimization of grammars for efficient communication. *Proceedings of the National Academy of Sciences*, *117*(5), 2347–2353.

Hawkins, J. A. (2004). *Complexity and efficiency in grammars*. Oxford: Oxford University Press.

Heine, B., & Kuteva, T. (2007). *The genesis of grammar: A reconstruction* (Vol. 9). Oxford University Press.

Kaufman, D. (2009). Austronesian nominalism and its consequences: A tagalog case study. *Theoretical Linguistics*, *35*(1), 1–49.

Smith, K., & Kirby, S. (2008). Cultural evolution: implications for understanding the human language faculty and its evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *363*(1509), 3591–3603.

Temperley, D., & Gildea, D. (2018). Minimizing syntactic dependency lengths: Typological/cognitive universal? *Annual Review of Linguistics*, *4*, 67–80.

Thompson, B., Kirby, S., & Smith, K. (2016). Culture shapes the evolution of cognition. *Proceedings of the National Academy of Sciences*, *113*(16), 4530–4535.

Travis, L. (1984). *Parameters and effects of word order variation*. Ph.D. dissertation, MIT.

Whitman, J. (2008). The classification of constituent order generalizations and diachronic explanation. In J. Good (Ed.), *Linguistic universals and language change* (p. 233-252). Oxford: Oxford University Press.