Regularisation, Systematicity and Naturalness in a Silent Gesture Learning Task

Yasamin Motamedi (Yasamin.Motamedi@ed.ac.uk)
Centre for Language Evolution, University of Edinburgh, 3 Charles Street, Edinburgh EH8 9AD, UK
Lucie Wolters (lwolters@exseed.ed.ac.uk)
Centre for Language Evolution, University of Edinburgh, 3 Charles Street, Edinburgh EH8 9AD, UK
Danielle Naegeli (d.s.u.naegeli@tilburguniversity.edu)
DCC, Tilburg University, Warandelaan 2, 5037 AB Tilburg, The Netherlands
Marieke Schouwstra (m.schouwstra@uva.nl)
ILLC, University of Amsterdam, Spuistraat 134, 1012 VB Amsterdam, The Netherlands
Simon Kirby (Simon.Kirby@ed.ac.uk)
Centre for Language Evolution, University of Edinburgh, 3 Charles Street, Edinburgh EH8 9AD, UK

Abstract

Typological analysis of the world's language shows that, of the 6 possible basic word orders, SOV and SVO orders are predominant, a preference supported by experimental studies in which participants improvise gestures to describe events. Silent gesture studies have also provided evidence for natural ordering patterns, where SOV and SVO orders are used selectively depending on the semantics of the event, a finding recently supported by data from natural sign languages. We present an artificial language learning task using gesture to ask to what extent preferences for natural ordering patterns, in addition to biases for regular languages, are at play during learning in the manual modality.

Keywords: silent gesture; constituent order; learning; regularization

Introduction

Languages can order the 3 basic clause constituents, subject (S), object (O) and verb (V), in 6 possible ways. However, these 6 possibilities are not uniformly distributed cross-linguistically; SVO and SOV orders comprise the basic word order for the considerable majority of the world's languages (Dryer, 2013; Napoli & Sutton-Spence, 2014).

A body of research using the silent gesture paradigm, where hearing participants produce gestures to communicate events without speech, have confirmed this preference experimentally, demonstrating that participants from different linguistic backgrounds produce SOV and SVO orders most frequently (Gibson et al., 2013; Goldin-Meadow et al., 2008; Hall et al., 2013; Meir et al., 2014). In particular, Schouwstra and de Swart (2014) investigated participants' ordering preferences for different types of events, suggesting that SOV and SVO order represent natural ordering for extensional and intensional events, respectively. Extensional events involve an action where a direct object is manipulated, usually involving movement through space (e.g. *throw*, *carry*), and where the direct object exists independently from the event itself. In contrast, intensional events are such that the meaning of the arguments, especially the direct object, are interpreted in relation to the event itself. That is, the object does not exist independently of the event, such as with creation events like *bake, paint* and *think*.

Under this account, ordering preferences are conditioned on the semantics of the event. For intensional events, where the existence of the direct object is dependent on the action denoted by the verb, SVO orders are produced most frequently, with the direct object following the verb. For extensional events, SOV orders are preferred. This finding has been demonstrated in silent gesture production studies with speakers of different language backgrounds (Schouwstra & de Swart, 2014), and for the comprehension of ordered gesture sequences as well as in gesture production (Schouwstra et al., 2019). However, these studies focus only on improvisation — i.e. the spontaneous creation of communicative behaviour in the absence of conventions. In a study using a similar semantic distinction between events, Christiansen et al. (2016) tested ordering preferences in gesture sequences produced by pairs of interacting participants in a communication task, demonstrating that natural preferences held after several rounds of communication. More recently, evidence from two sign languages, Brazilian Sign Language (Libras) and Nicaraguan Sign Language (NSL) suggest that semantically-conditioned word order preferences are present in natural languages (Flaherty et al., 2018; Napoli et al., 2017). These systems are, of course, not improvised, but conventional linguistic systems that are the product of transmission to new learners over time. The question remains, however, how these semantic distinctions persist or arise in natural language, and whether the bias for natural ordering patterns seen in improvisation studies applies during learning.

While improvisational studies such as those described above are a valuable tool for probing which cognitive biases are at play in the absence of any conventions, they are not the only methodological tool to investigate cognitive biases during different stages of communication. For example, artificial language learning (ALL) paradigms, in which participants are trained on artificially constructed 'languages' before being tested on what they have learnt, have been used widely to test hypotheses related to the cognitive biases that underpin language learning under different constraints (Culbertson et al., 2012; Fedzechkina et al., 2012; Ferdinand et al., 2019; Hudson Kam & Newport, 2005). Culbertson et al. (2012) used an ALL paradigm to investigate biases for noun phrase ordering patterns. As with ordering of basic

constituents, noun phrase ordering patterns are not uniformly distributed across languages, but harmonic patterns, where the position of the modifier is consistent relative to the noun, are more common than non-harmonic orders. Culbertson and colleagues trained participants on different input 'languages' with different proportions of each ordering pattern, showing that participants' ability to learn (and further regularise) the more systematic harmonic patterns was greater than learning of non-harmonic patterns, and thus offers evidence for an individual-level learning bias shaping language typology. Ferdinand et al. (2019) used an ALL task alongside a domain-general learning task to demonstrate biases for regularisation in learning. Their results showed that regularisation occurred in both tasks, though more so in the language learning task, suggesting both domain-general and domain-specific sources for regularisation biases during learning.

While silent gesture studies focussing on word order have rarely looked beyond the improvisation stage, ALL experiments have largely focused on written or spoken language. The use of the manual modality in silent gesture experiments, as a potentially linguistic modality used by participants for whom it is *not* linguistic (i.e. spoken language users), can further reduce the influence participants' existing linguistic knowledge will bring to bear on the task. Furthermore, the question remains whether the same learning biases operate across different linguistic modalities.

Here, we apply an ALL paradigm similar to the design of the linguistic task used by Ferdinand et al (2019) to the manual modality, training participants on gesture sequences shown at different frequencies before assessing which gesture sequences they select in a testing phase. Participants in our study were shown gesture sequences denoting either an extensional or an intensional event — those sequences were ordered either with SOV or SVO order, and we manipulated the frequency with which they saw each order with each event type during training. In a forced-choice testing stage, we presented participants with each event type repeatedly over several trials and asked them to select a gesture sequence in a two forced-choice task, between either an SOV- or SVO-ordered sequence.

We ask firstly whether participants learn the ordering patterns of silent gesture sequences, based on the frequency with which they see each order-event mapping during the training stage. We also ask how participants change the input they learn from. Do participants regularise the input frequencies they are trained on, or do they probability match? In order words, are their output behaviours less variable than training, and if so, how does this reduction in variability compare to that seen in other modalities? Do participants systematise their input in favour of a single order? Is there a tendency for participants to produce output behaviours that treat the two event types more similarly than their training input? Finally, we ask whether participants reproduce or overproduce the natural mappings seen in training, such that SOV order corresponds to extensional events and SVO order to intensional events. That is, do we replicate in a learning task the naturalness bias seen in improvisation tasks?

Methods

Participants 200 participants were recruited for an online experiment using the online crowdsourcing platform Prolific (www.prolific.co). Participants were prescreened to have English as their first language, and to have not participated in any of our previous silent gesture experiments. We excluded 1 participant due to an experiment error. The monetary compensation per participant was equivalent to £12.35 per hour.

Materials The experiment ran as a web app in the participant's web browser, using the JSpsych javascript library. Participants were shown 2 line drawings depicting different events (shown in figure 1) and 4 gesture videos throughout the experiment. The line drawings represented the extensional event 'nun throws ukulele', and the intensional event 'nun thinks of ukulele'. Both pictures corresponded to two videos displaying a member of the research team using 3 gestures (and no speech) to convey the information in the event, one using SOV word order (*nun-ukulele-throw/think*), the other SVO word order *(nun-throw/think-ukulele*). The distinct constituents conveyed in the videos were each gestured with a single unique gesture. All videos were 4.5 seconds long and can be accessed via our OSF page.

Figure 1. Event pictures used in the study, showing the extensional event *nun-throws-ukelele* (left) and the intensional event *nun-thinks of-ukelele* (right).

Procedure The experiment consisted of a training phase and a production phase. At the start of the training phase, participants were instructed to sit back and watch carefully. In each training trial, an event picture was displayed for 1000ms and then a gesture video corresponding to the event was displayed below the picture for the length of the video. In each production trial, an event picture was displayed with both corresponding gesture videos. The videos looped simultaneously and were displayed side-by-side below the event picture. The locations (left or right) of the videos were randomised per trial per participant. Participants were instructed to click on the gesture sequences "like they saw in the first part of the experiment". Both the training and production phase consisted of 20 trials showing 10 extensional and 10 intensional events in an order randomized per experiment phase per participant. Participants were not informed in advance about the number of trials in the study.

Design The experiment had a between-subjects design, in which participants were randomly assigned to one of four

training conditions, which varied according to the number of times participants saw videos with SOV or SVO order in the training phase. Each condition featured a majority order and a minority order, for each event type. The majority gesture video was displayed in 7/10 trials, and the minority gesture video in 3/10 trials. The order of presentation of the 20 trials was randomised per participant. In the *natural* condition (N=50), the majority orders reflected natural semantic mappings, following Schouwstra and de Swart (2014) – participants saw the extensional events gestured with an SOV majority, and the intensional events with an SVO majority order. In the *unnatural* condition (N=48), the natural semantic mappings were inverted, such that participants saw the extensional events gestured with an SVO majority, and the intensional events with an SOV majority order. In the *majority SVO* condition (N=49), participants saw both event types gestured with an SVO majority order, and in the *majority SOV* (N=52), condition participants saw both event types gestured with an SOV majority order. Table 1 summarises the input orders for each condition. The experimental design and analysis plan were pre-registered on the Open Science Framework prior to data collection[1].

Table 1 Input proportions of SVO and SOV orders for each event type in each condition

Condition	Extensional events	Intensional events
natural	70% SOV, 30% SVO	30% SOV, 70% SVO
unnatural	30% SOV, 70% SVO	70% SOV, 30% SVO
majority SVO	30% SOV, 70% SVO	30% SOV, 70% SVO
majority SOV	70% SOV, 30% SVO	70% SOV, 30% SVO

Results

Learning

Firstly, we analysed whether participants' responses show evidence of learning from the input they receive. We assessed whether participants' selections at each trial matched the majority order they saw in training for that event. Figure 2 illustrates the proportion of trials that matched the majority order. We analysed our data using a mixed effects logistic regression, implemented with R (R Core Team, 2013) and lme4 (Bates et al., 2015). Our model included fixed effects of condition, event type (extensional/intensional) and their interaction, with the outcome variable being a binary variable noting whether or not participants matched the majority order at each trial. We included a by-participant random intercept, with a random slope of event type. Both condition and event type were deviation coded. Model comparison revealed that a simpler model, without event type, represented the best fit to the data ($\chi = 29.61$, $p < 0.001$).

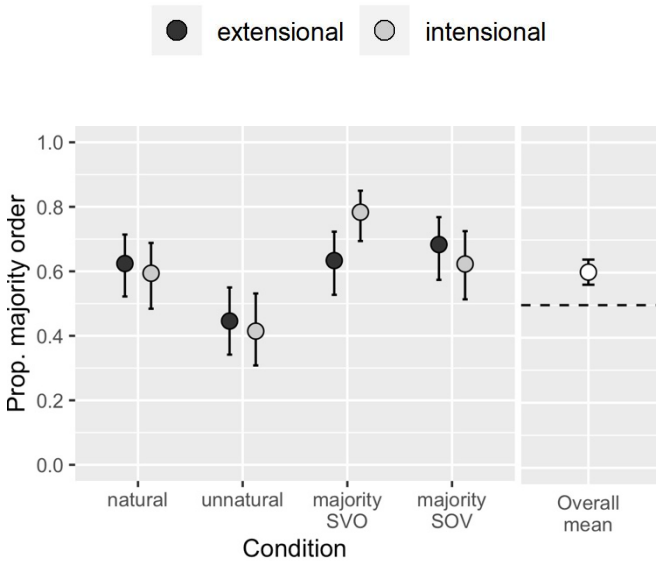

Figure 2: Plot showing the proportion of test trials in which participants select the input majority order, shown for each condition and each event type (shaded points), as well as the overall mean across conditions (right). Error bars represent bootstrapped 95% confidence intervals. Participants' selection of the majority trained order indicates learning, except in the unnatural condition; overall, participants select the majority order more frequently than would be expected by chance.

The model revealed a significant intercept ($\beta = 0.83$, SE = 0.16, $z = 5.32$, $p < 0.001$), suggesting that, on average, participants across conditions select the majority order more than would be expected by chance. The model also revealed effects of condition for the unnatural condition ($\beta = -1.45$, SE = 0.27, $z = -5.32$, $p < 0.001$), such that participants produce the majority order *less* often, and in the majority SVO condition ($\beta = 0.95$, SE = 0.28, $z = 3.39$, $p < 0.001$), such that participants produce the majority order more often in this condition. We relevelled the model to extract the coefficients for the natural condition; we find no condition effects for the natural or the majority SOV conditions.

Systematisation

We define systematicity as the case where meanings across the system are treated the same way – i.e. the most systematic ordering preference would use the same constituent order for all event descriptions. In this way, we can measure an

[1] Pre-registration and all analysis files available at https://osf.io/4wnjv/

increase in systematicity as a reduction in the variation across a system.

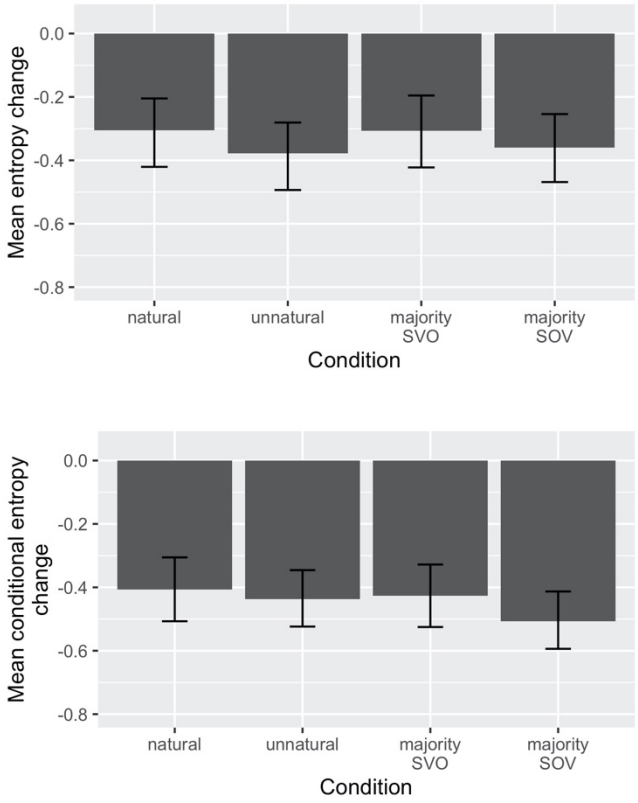

Figure 3: Plots showing mean change between training input and participant output for entropy (top) and conditional entropy (bottom), relating to systematisation and regularisation respectively. Error bars represent bootstrapped 95% confidence intervals. Both systematisation and regularisation arise from learning in this task.

Following Ferdinand et al. (2019), we quantify the variation in a system using Shannon entropy across meanings in each condition[2], where the entropy (*H)* of a system is given as:

$$H(V) = -\sum_{v_i \in V} p(v_i) log_2\, p(v_i)$$

where, *V* is the set of variants (here SOV and SVO orders). For example, the natural and unnatural conditions both have an entropy value of 1 for the input that participants see in training because each order occurs in 50% of trials. We calculate the change in entropy between the input participants receive and the output they produce, to assess whether participants are biased to produce outputs that are more systematic than the input they receive.

Figure 3 (top) illustrates the mean entropy change in each condition and the confidence intervals around each condition mean. As the distribution of entropy values in our study is non-normal, our data do not meet the assumptions for the linear models that are often applied to entropy data. Instead, we calculated bootstrapped confidence intervals around the mean entropy change in each condition (shown in table 2), using the 'boot' package in R, generating 10,000 samples. We extracted 95% confidence intervals using the accelerated bias-corrected method (BCa), as recommended by Puth et al. (2015). We also calculated 95% confidence intervals around the differences in condition means using the same methods (table 3). None of the confidence intervals around condition means contain zero, suggesting a reliable reduction in entropy for all conditions. Similarly, confidence intervals around differences between conditions all contain zero, suggesting that we cannot reliably identify differences between conditions.

Table 2. Mean entropy change and 95% bootstrapped confidence intervals around the mean for each condition.

Condition	$\bar{x}$	Lower CI	Upper CI
natural	-0.31	-0.43	-0.21
unnatural	-0.38	-0.49	-0.28
majority SVO	-0.31	-0.42	-0.20
majority SOV	-0.36	-0.46	-0.26

Table 3. Mean difference in entropy change and 95% bootstrapped confidence intervals around the mean difference between conditions.

Conditions	$\bar{x}_a - \bar{x}_b$	Lower CI	Upper CI
nat – unnat	0.07	-0.08	0.22
nat – SVO	0.001	-0.15	0.16
nat – SOV	0.05	-0.10	0.20
unnat – SVO	-0.07	-0.22	0.08
unnat – SOV	-0.02	-0.17	0.13
SVO – SOV	0.05	-0.10	0.21

Regularisation

In contrast to systematisation, which reflects an increase in similarity in word order *across* categories, regularisation refers to a decrease in variability *within* a category. Accordingly, we measure regularisation as a drop in conditional entropy, which takes into account the probability of variants appearing in particular contexts. The conditional entropy of a system is given as

$$H(V|C) = -\sum_{c_j \in C} p(c_j) \sum_{v_i \in V} p(v_i|c_j) log_2\, p(v_i|c_j)$$

[2] Note that this measure differs from the planned measure described in our pre-registration, which operationalised systematicity as the extent to which participants produced either of the two orders as a majority order. All analysis files, including the pre-registered analysis, can be found at https://osf.io/4wnjv/ .

where V is the set of variants (here SOV and SVO orders) and C is the set of contexts the variants appear in (here, our event types). We calculated the change in conditional entropy between the input participants received ($H(V|C) = 0.88$ across all conditions) and the output they produced (illustrated in the bottom panel of figure 3). We calculated bootstrapped 95% confidence intervals around the mean for each condition (table 4), using the methods described for our systematisation measure above, as well as around the differences in condition means (table 5). Our findings suggest that conditional entropy reduces across conditions; we do not find reliable differences in conditional entropy change between conditions.

Table 4. Mean conditional entropy change and 95% bootstrapped confidence intervals around the mean for each condition.

Condition	$\bar{x}$	Lower CI	Upper CI
natural	-0.41	-0.51	-0.31
unnatural	-0.44	-0.52	-0.35
majority SVO	-0.43	-0.52	-0.32
majority SOV	-0.51	-0.59	-0.42

Table 5. Mean difference in conditional entropy change and 95% bootstrapped confidence intervals around the mean difference between conditions.

Conditions	$\bar{x}_a - \bar{x}_b$	Lower CI	Upper CI
nat – unnat	0.03	-0.10	0.17
nat – SVO	0.02	-0.12	0.16
nat – SOV	0.10	-0.04	0.23
unnat – SVO	-0.01	-0.14	0.12
unnat – SOV	0.07	-0.06	0.20
SVO – SOV	0.08	-0.05	0.21

Naturalness

Finally, we analysed whether participants across conditions show a bias in favour of natural ordering patterns. Figure 4 illustrates the proportion of trials where the selected order matches the expected natural order (extensional = SOV, intensional = SVO). We ran a logistic mixed effects model on the binary variable noting whether selected order matched natural order, with a model structure identical to that described above for learning. Model comparison revealed that the full model (with the interaction term) represented the best fit in this case ($\chi = 34.6$, $p < 0.001$).

Analysis of the model results shows a significant intercept ($\beta = 0.62$, SE = 0.11, 5.52, $p < 0.001$), indicating that, on average, participants selected the natural order more often than we would expect by chance (note that collapsed across conditions, natural order occurred 50% of the time). The model revealed no significant main effects, but did show significant interactions between event type and the two majority order conditions. The interaction between event type and majority SVO condition ($\beta = 3.64$, SE = 0.71, $z = 5.16$, $p < 0.001$), suggests that natural order is used more often for intensional events, where natural order is consistent with the majority. Conversely, the interaction between event type and the majority SOV condition (β = -2.99, 0.68, -4.40, $p < 0.001$), suggests that the preference for natural order is lower for intensional events, where it conflicts with the input majority.

Figure 4: Plot showing the proportion of test trials in which

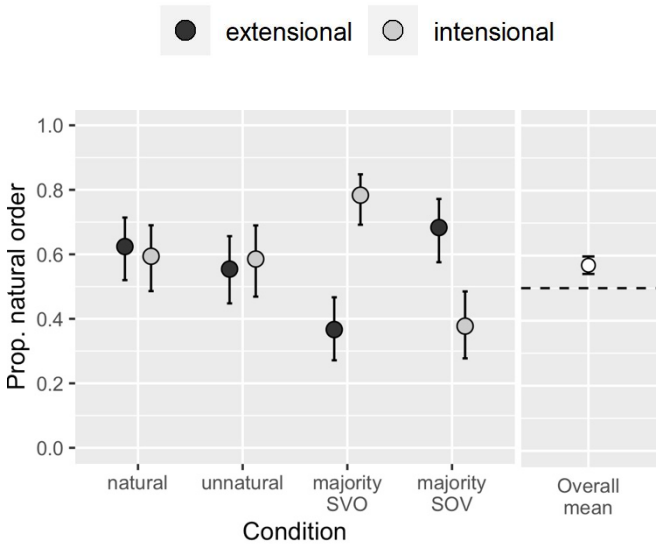

participants select the natural order, shown for each condition and each event type (shaded points), as well as the overall mean (right). Error bars represent bootstrapped 95% confidence intervals. Overall, participants select the natural order more frequently than would be expected by chance.

Discussion

We have presented a study that investigates whether biases present during the improvisation of gestural signals persist during learning, further extending ALL research on regularisation to the manual modality.

Firstly, we have shown that participants are able to learn from gestural input in an online ALL experiment, with a similar design to studies focussing on written and spoken language. Participants select orders in testing that closely reflect the majority orders they received in their training input, with the exception of the unnatural condition. Notably, we see that participants in the unnatural condition, where the input orders should be in conflict with a bias for naturalness and a bias for systematisation, do not reproduce very well the ordering patterns seen in training. We also find overproduction of the majority order for the majority SVO condition, which may reflect biases from our participants' native language, English – however, it is not the case that across conditions participants overproduce SVO, only when SVO occurs in the majority in the input.

We also find that participants' outputs differ from their training input in several ways. Firstly, participants across conditions show a tendency to systematise from the input, reflected in the change in Shannon entropy between the training input and participants' output, such that ordering

preferences across event types become more similar. This possibly reflects a general bias for simplicity (Culbertson & Kirby, 2016), where the simplest system would be one in which all events are represented with a single order (i.e. either SOV or SVO).

Consistent with Ferdinand et al. (2019), we find that participant's outputs show regularisation of the training input, evidenced through a reduction in conditional entropy between input and output. Importantly, the extent of this regularisation is consistent with the linguistic task reported by Ferdinand et al. (2019), showing substantially more regularisation than their non-linguistic task (regularising on average 0.36 bits in the linguistic task compared to 0.17 bits in the non-linguistic task). This finding suggests that participants learn in a silent gesture task in a similar way to previous ALL tasks, focussed on learning in spoken languages. If these typical artificial language tasks are analogous to learning in spoken language, we suggest that our findings point to regularisation processes being highly similar across both spoken and signed languages.

Finally, a key question of this work was to understand whether the same bias for natural ordering patterns that we see in improvisation tasks (Christensen et al., 2016; Goldin-Meadow et al., 2008; Schouwstra & de Swart, 2014) is also present during learning. Our results suggest that it is. Participants reproduce and even extend natural ordering patterns in their output. In particular, in the unnatural condition, where the input orders appear to be dispreferred (as evidenced from the learning results), participants' output actually shows a similar proportion of natural orders to the natural condition, despite the input being the inverse. This overall preference for natural orders may explain why we can see similar ordering patterns, where word/sign order is conditioned on event, type in natural sign languages such as Libras and NSL (Flaherty et al., 2018; Napoli et al., 2017), and not just in improvised gesture.

One open question is to what extent we expect this semantically conditioned ordering pattern to persist, given that such a distinction has not been found in the majority of the world's languages. Indeed, systematic ordering patterns, where languages tend to use a single word order independent of event type, appears to be the norm. To date, semantically conditioned ordering patterns of the kind studied here have only been reported in two sign languages, and a number of factors may make the appearance of these orders in sign versus spoken languages more likely, such as age of the language, community structure (Meir et al., 2005) and iconicity (Christensen et al., 2016; Hall et al., 2013). Future work should look at the factors that influence competition between naturally conditioned and systematic, *un*conditioned order as a language continues to evolve.

References

Bates, D., Bolker, B., Machler, M., & Walker, S. C. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, *67*(1), 1–48.

Christensen, P., Fusaroli, R., & Tylén, K. (2016). Environmental constraints shaping constituent order in emerging communication systems: Structural iconicity, interactive alignment and conventionalization. *Cognition*, *146*, 67–80.

Culbertson, J., & Kirby, S. (2016). Simplicity and Specificity in Language: Domain-General Biases Have Domain-Specific Effects. *Frontiers in Psychology*, *6*, 1–11.

Culbertson, J., Smolensky, P., & Legendre, G. (2012). Learning biases predict a word order universal. *Cognition*, *122*(3), 306–329.

Dryer, M. S. (2013). Order of Subject, Object and Verb. In M. S. Dryer & M. Haspelmath (Eds.), *The World Atlas of Language Structures Online*. Max Planck Institute for Evolutionary Anthropology. https://wals.info/chapter/81

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

Ferdinand, V., Kirby, S., & Smith, K. (2019). The cognitive roots of regularization in language. *Cognition*, *184*, 53–68.

Flaherty, M., Schouwstra, M., & Goldin-Meadow, S. (2018). Do we see word order patterns from silent gesture studies in a new natural language? In C. Cuskley, M. Flaherty, H. Little, L. McCrohon, A. Ravignani, & T. Verhoef (Eds.), *The Evolution of Language: Proceedings of the 12th International Conference (EVOLANGXII)*. NCU Press.

Gibson, E., Piantadosi, S. T., Brink, K., Bergen, L., Lim, E., & Saxe, R. (2013). A Noisy-Channel Account of Crosslinguistic Word-Order Variation. *Psychological Science*, *24*(7), 1079–1088.

Goldin-Meadow, S., So, W. C., Ozyürek, A., & Mylander, C. (2008). The natural order of events: How speakers of different languages represent events nonverbally. *PNAS*, *105*(27), 9163–9168.

Hall, M. L., Mayberry, R. I., & Ferreira, V. S. (2013). Cognitive constraints on constituent order: Evidence from elicited pantomime. *Cognition*, *129*(1), 1–17.

Hudson Kam, C. L., & Newport, E. (2005). Regularizing Unpredictable Variation: The Roles of Adult and Child Learners in Language Formation and Change. *Language Learning and Development*, *1*(2), 151–195.

Meir, Irit, Israel Assaf, Wendy Sandler, Carol Padden, and Mark Aronoff. 'The Influence of Community on Language Structure'. *Linguistic Variation* 12, no. 2 (2005): 247–91.

Meir, I., Aronoff, M., Börstell, C., Hwang, S., Ilkbasaran, D., Kastner, I., Lifshitz, A., Basat, B., Padden, C., & Sandler, W. (2014). The effect of being human and the basis of grammatical word order. *Cognition*, *158*(April), 1–40.

Napoli, D. J., Spence, R. S., & Quadros, R. M. de. (2017). Influence of predicate sense on word order in sign languages: Intensional and extensional verbs. *Language*, *93*(3), 641–670.

Napoli, D. J., & Sutton-Spence, R. (2014). Order of the major constituents in sign languages: Implications for all language. *Frontiers in Psychology*, *5*, 1–18.

Puth, M.-T., Neuhäuser, M., & Ruxton, G. D. (2015). On the variety of methods for calculating confidence intervals by bootstrapping. *Journal of Animal Ecology*, *84*(4), 892–897.

R Core Team. (2013). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. http://www.R-project.org/

Schouwstra, M., & de Swart, H. (2014). The semantic origins of word order. *Cognition*, *131*(3), 431–436.

Schouwstra, M., Swart, H. de, & Thompson, B. (2019). Interpreting Silent Gesture: Cognitive Biases and Rational Inference in Emerging Language Systems. *Cognitive Science*, *43*(7), e12732.