Cognition 185 (2019) 49–61

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Original Articles

Acquisition and processing of an artificial mini-language combining semantic and syntactic elements

Fosca Al Roumi[a,*], Dror Dotan[a,b], Tianmin Yang[c], Liping Wang[c], Stanislas Dehaene[a,d]

[a] Cognitive Neuroimaging Unit, CEA DRF/I2BM, INSERM, Université Paris-Sud, Université Paris-Saclay, NeuroSpin Center, France
[b] Mathematical Thinking Lab, School of Education and the Sagol School of Neuroscience, Tel Aviv University, Israel
[c] Institute of Neuroscience, Key Laboratory of Primate Neurobiology, Shanghai Institutes for Biological Sciences, Chinese Academy of Sciences, Shanghai 200031, China
[d] Collège de France, Paris, France

ARTICLE INFO

Keywords:
Syntax
Finger tracking
Sequence learning
Artificial grammar

ABSTRACT

Most artificial grammar tasks require the learning of sequences devoid of meaning. Here, we introduce a learning task that allows studying the acquisition and processing of a mini-language of arithmetic with both syntactic and semantic components. In this language, symbols have values that predict the probability of being rewarded for a right or left response. Novel to our paradigm is the presence of a syntactic operator which changes the sign of the subsequent value. By continuously tracking finger movement as participants decided whether to press left or right, we revealed the successive cognitive stages associated with the sequential processing of the semantic and syntactic elements of this mini-language. All participants were able to understand the semantic component, but only half of them learned the rule associated with the syntactic operator. Our results provide an encouraging first step in elucidating the way in which humans acquire non-verbal syntactic structures and show how the finger tracking methodology can shed light on real-time artificial language processing.

1. Introduction

Infant and adult humans are able to extract and compress the information present in their rich environment. From a small number of samples, they can identify the abstract hierarchical rules that govern the structure of the incoming information (Amalric et al., 2017). This ability lies at the heart of the human ability to acquire the syntactic and semantic structures of human language. It has therefore been proposed that humans, contrary to other animals, possess the specific ability to acquire recursive rules (Fitch, 2004; Hauser, 2002; ten Cate & Okanoya, 2012).

To study this ability in the laboratory, researchers have used mini-languages involving minimal lexicons and grammatical rules. This artificial grammar approach (Reber, 1969) has led to important advances in understanding the mechanisms of rule acquisition and generalization in infants. For instance, it showed that infants can use transition probabilities to identify the chunks that form words and the abstract patterns that govern sequences of syllables or tones (Marcus et al., 1999; Saffran, Johnson, Aslin, & Newport, 1999). The comparison of such data with formal models of children's learning (Frank & Tenenbaum, 2011) is currently being used to clarify the statistical learning mechanism and the language primitives that drive human language learning.

To understand which rule learning abilities allow humans to acquire and use complex languages, many studies compare them with their closest phylogenetic relatives, non-human primates, in similar tasks of sequence processing and artificial grammar learning (Wilson, Marslen-Wilson, & Petkov, 2017). Several of the mechanisms involved in sequence processing seem to be shared by human and non-human primates and could reflect domain-general abilities contributing to rule learning and language acquisition in humans. For instance, a recent study suggests that primates may process sequential relationships in a way strikingly similar to human infants (Milne et al., 2016). Moreover, non-human primates have been shown to extract the statistical properties of syllable transition probabilities in a continuous speech stream in a manner comparable to humans (Hauser, Newport, & Aslin, 2001). Furthermore, when processing the order relationships between non-sense words in rule-based sequences, frontal brain regions activate similarly in macaques and humans (Wilson, Kikuchi, et al., 2015).

However, primates' grammatical abilities seem to diverge from those of humans whenever embedding and recursion are needed to process the sequences (Penn, Holyoak, & Povinelli, 2008). For instance, human and non-human primates have been shown to master simple grammars devoid of embedding. Yet only humans succeeded to do so

* Corresponding author at: CEA.DSV.I2BM.NeuroSpin – Cognitive Neuroimaging Unit, Centre d'études de Saclay, 91191 Gif-sur-Yvette, France.
E-mail address: fosca.alroumi@gmail.com (F. Al Roumi).

https://doi.org/10.1016/j.cognition.2018.11.006
Received 7 January 2018; Received in revised form 16 November 2018; Accepted 19 November 2018

when the grammar contained center-embedded sequences (Fitch, 2004). Furthermore, both macaques and humans showed sensitivity to adjacent relationships, but significant sensitivity to the non-adjacent relationships of a mixed-complexity artificial grammar was found only in a subgroup of humans (Wilson, Smith, & Petkov, 2015). We note that both humans and monkeys failed to learn the nonadjacent dependencies using a similar paradigm but with different stimuli (Milne, Petkov, & Wilson, 2018). Other works seem to show that non-human primates can process nonadjacent dependencies in some artificial grammar learning tasks (e.g. see Newport, Hauser, Spaepen, & Aslin 2004; Ravignani, Sonnweber, Stobbe, & Fitch, 2013; Sonnweber, Ravignani, & Fitch, 2015; Milne et al., 2016). Finally, a recent fMRI study suggested that both humans and macaques are capable of representing the number and sequence patterns underlying simple auditory sequences, but that only humans possess an integrated representation of those two features in the inferior frontal gyrus (Wang, Uhrig, Jarraya, & Dehaene, 2015).

One issue with such artificial grammar studies is that they typically require only the learning of the surface structure of sequences (their 'syntax', i.e. the rules governing the valid or invalid arrangement of their elements), but the sequences do not convey any meaning. Indeed, in these artificial grammars, no semantic component is present and only a minimal form of syntax, often reducible to a finite-state automaton, is needed in order to detect grammatical violations (for rare exceptions, see Friederici, Steinhauer, & Pfeifer, 2002; Moeser & Olson, 1974; Mueller, Hahne, Fujii, & Friederici, 2005). As the primary role of syntax is to convey meaning, these studies may fail to engage some crucial mechanisms involved in semantic processing and language acquisition in humans. Our goal in the present work was to fill this gap. To do so, we designed a simple language with elementary syntactic and semantic features and studied its learnability in human adults. Such a study is a first step prior to comparing its learnability by human and non-human primates.

We designed an artificial mini-language reminiscent of elementary symbolic arithmetic, and studied whether it could be learned by trial-and-error, in the absence of explicit instructions. Each 'sentence' of the language was a sequence of semantic and syntactic elements, represented by visual symbols. The first part of the experiment was very similar to Yang and Shadlen (2007): on each trial, participants saw a sequence of 1, 2 or 4 symbols (sampled randomly out of a set of six symbols) and had to point to a left or right response button. The rewarded response was determined by merging all symbols: each symbol was associated with a certain weight of evidence (hereby*WOE*) favoring a decision to the right or to the left, and the probability of the rewarded response being left or right was determined by summing the *WOE*s of all symbols in the sequence (Fig. 1a, b). For instance, if symbol A strongly favored a rightward response and symbol B slightly favored a leftward response, then participants had to compute that the sequence AB still slightly favored a rightward response (according to a precise quantitative equation, see Section 2.2). Note that the syntax here is essentially non-existent, because the semantics is order-independent: all symbol sequences are meaningful, and their meaning is simply the sum of the weights attached to each symbol, which is a commutative operation. Previous work by Yang and Shadlen (2007) demonstrated that even rhesus macaques could learn such a language: they acquired a lexicon of 10 symbols, learned their weights, and learned how to sum these weights for decisions – yet with training that required more than 130,000 trials. In the first part of this study, we will see that humans can learn this language too, within about an hour and ~1,000 trials.

The second part of this study aimed to examine how humans learn a language with both syntax and semantics. To this end, we added to the language above a syntactic operator, represented by another visual symbol. Unlike the other symbols, this symbol did not have a fixed value, but its function was to invert the sign of the *WOE* of the subsequent symbol. For instance, if this symbol is denoted by the ⋆ sign and symbols A and B are defined as above (A → right, B → left), then the sequence of the 3 symbols A ⋆ B strongly favored a rightward response, whereas the sequence B ⋆ A strongly favored a leftward response. The analysis of the participants' responses to this order-dependent (non-commutative) mini-language allowed us to determine whether they were able to learn the meaning of the syntactic symbol. A participant who fails to learn the symbol may come up with an alternative interpretation for it – e.g. assign a *WOE* to the syntactic symbol just like to any other semantic symbol. Because this mini-language is an extension of the language considered in Yang and Shadlen (2007), which is learnable by monkeys, we reasoned that the extended language, if easily learnable by humans, could pave the way to a future investigation of syntactic abilities in non-human primates.

The present work was additionally designed to show that a purely behavioral measure, namely continuous finger tracking, can be sensitive to the covert mechanisms that are engaged in the sequential processing of symbols. On each trial, we recorded the finger movement while participants saw symbol sequences and pointed to one of two response buttons on a touchscreen. Crucially, the participants had to move their finger continuously without ever stopping. We also enforced a minimal finger speed in order to encourage intuitive decisions and dissuade explicit calculations. Previous research has repeatedly shown that intermediate points in the pointing trajectory reflect the real-time unfolding of internal decision processes (Berthier, 1996; Erb, Moher, Sobel, & Song, 2016; Friedman, Brown, & Finkbeiner, 2013; Pinheiro-Chagas, Dotan, Piazza, & Dehaene, 2017; Resulaj, Kiani, Wolpert, & Shadlen, 2009; Song & Nakayama, 2009). We therefore hypothesized that the continuous finger tracking would provide insights into the online processing of the symbol sequence, without requiring invasive neurophysiological recordings of single unit activity (Yang & Shadlen, 2007), which are typically not available in human subjects.

2. Methods

2.1. Participants

The participants were 30 French right-handed adults (this number was determined based on past experience with similar paradigms Dotan, Meyniel, & Dehaene, 2018). We excluded 3 participants who performed poorly in the length-1 sequences of Day 1 (2 standard deviations below the mean). 27 participants (16 females, mean age = 23.7 years, SD = 2.8 years) took part to the first day of the experiment, and 25 returned on the successive day for the second part. All participants gave informed consent prior to participating. The study was approved by the local ethics committee.

2.2. Stimuli

On each trial, the participants saw 1, 2 or 4 symbols, presented sequentially, and were instructed to point to one of two response buttons (Fig. 1). In most experimental blocks, one of the buttons was rewarded by presenting correct/incorrect feedback at the end of the trial. As explained above, each sequence of symbols was associated with a certain probability that the rewarded response button would be the one on the right (hereby P^{right}). There were 6 "content symbols", each associated with a certain weight of evidence (hereby *WOE*), and the sum of *WOE*, ΣWOE, determined $P^{right} = (1 + 10^{-\sum WOE})^{-1}$ (Fig. 1a). The sequences of stimuli were pseudo randomized, such that each "content symbol" appeared the same number of times at each position. A 7th symbol, denoted ⋆, was not associated with a *WOE*, but had a syntactic role: it inverted the sign of the *WOE* of the subsequent symbol (but left the *WOE* of the other symbols unchanged). Namely, the ⋆ operator acted locally, affecting only the symbol following it. Throughout this article, ▲ denotes any content symbol and ⋆ denotes the syntactic operator. Specific content symbols are denoted according to their *WOE*, as $<\ <\ <, <\ <, <, >, >\ >, >\ >\ >$.

An ideal observer should choose the response button corresponding

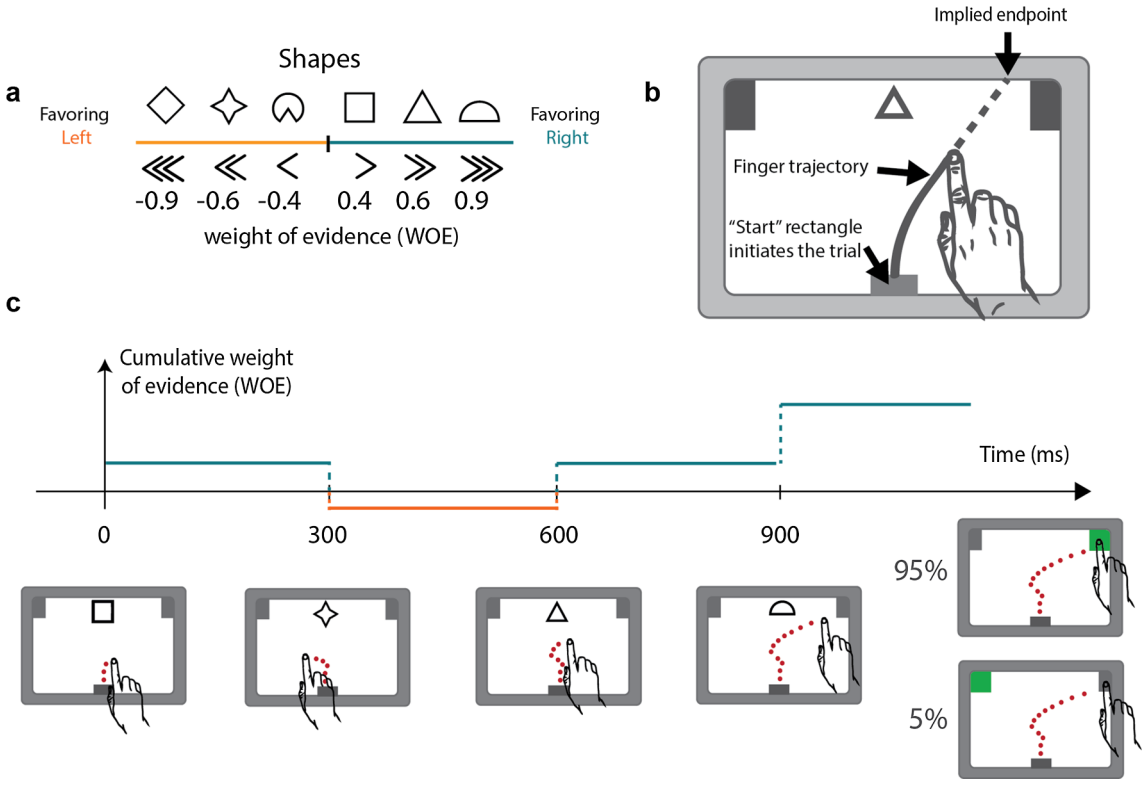

Fig. 1. Experimental paradigm. On each trial, one, two or four symbols were presented sequentially on an iPad. Participants dragged their finger from the starting point to either a left or a right response button. (a) Meaning of the symbols. Each of the 6 symbols was assigned a certain weight-of-evidence (*WOE*) that determined the probabilities of reward. The symbol-WOE mapping was fixed for each subject. The 6 symbols are represented in the text according to their *WOE* as $<\ <\ <$, $<\ <$, $<$, $>$, $>\ >$, and $>\ >\ >$. (b) Screen layout. Finger movement was continuously recorded. The onset of the first symbol was triggered by finger movement. After touching a response button, feedback was provided. (c) Sample trial flow. The probability that the right or left button would be rewarded was a sigmoid function of the sum of the weights of the symbols in the sequence. For the specific sequence of symbols shown here, the probability of obtaining a positive feedback for a right response is 95%.

Table 1
The experiment design. The experiment took place during two consecutive days. There were 4 blocks in each day, in increasing level of difficulty. In the first day, only semantic symbols were presented (each denoted here by ▲). In the second day, the sequences also included the syntactic operator ⋆.

	DAY 1			DAY 2		
	Sequences	**#Trials**	**Feedback**	**Sequences**	**#Trials**	**Feedback**
length-1	▲	480	Yes	▲	240	Yes
length-2	▲▲	360	Yes	⋆▲ ▲▲	360	Yes
length-4	▲▲▲▲	360	Yes	⋆▲▲▲ ▲⋆▲▲	480	Yes
		360	No	▲▲⋆▲ ▲▲▲▲	360	No

to the side that agrees with $sign(\sum WOE)$. Such a "maximizing strategy" does not guarantee a reward, i.e. a positive feedback, because even the highest ΣWOE did not reach a P^{right} value of 1.0 or 0.0, but it does maximize the chances of getting such feedback.

2.3. Trial flow

Touching the "start" rectangle (Fig. 1b) triggered the onset of a fixation dot at the symbols' intended location (top center, see Fig. 1b). When the participant started moving their finger upwards, the symbol sequence started appearing. Each symbol was presented for 100 ms, with a stimulus onset asynchrony (SOA) of 300 ms. As soon as the finger touched one of the response buttons, feedback was provided: the button that should have been selected flickered in green, and a pleasant/unpleasant sound was played. A trial was considered as failed if the participant lifted the finger before reaching the top of the screen, moved the finger backwards, started the trial with sideways (rather than upward) movement or moved too slowly (more than 3 s per trial or more than 1.5 s to reach the first third of the screen, with linear interpolation; except a grace period of the trial's first 300 ms). In such cases, the trial immediately terminated and was presented again later in the experiment.

2.4. Training

To get acquainted with the motor aspects of the task, the participants initially performed a similar task in which the stimulus was a single arrow pointing left or right, which unambiguously indicated the response button that would be rewarded.

2.5. Experiment stages

The experiment was organized in two sessions, held on two consecutive days. Each session lasted approximately 90 minutes and was organized in 4 parts (see Table 1). During the first day, only the 6 content symbols were presented; in the second day the ⋆ operator was added.

To examine whether the participants genuinely understood the syntactic operator ⋆, rather than just memorized specific symbol combinations, we tested their ability to generalize the function of the syntactic operator to untrained combinations. Thus, in all blocks with feedback, the ⋆ operator never occurred prior to the symbol $<\ <$ or $>\ >$. Such combinations of ⋆ followed by $<\ <$ or $>\ >$ were presented only during the last stage in day 2 (with length-4 trials), for which no feedback was provided.

2.6. Questionnaire

At the end of the second day, participants filled out a questionnaire where they were asked for the meaning of each of the 7 symbols, i.e. how each symbol contributed to their left/right decision. We used this questionnaire to divide the participants into two groups – those who managed to figure out explicitly the meaning of the syntactic operator ($G^{\star+}$ participants) and those who did not ($G^{\star-}$ participants).

2.7. Data encoding

The position of the finger was sampled at 60 Hz and resampled to 100 Hz using cubic spline interpolation. The x and y coordinates were then separately smoothed with a Gaussian filter with $\sigma = 20$ ms. At each time point, we calculated the implied endpoint (*iEP*) – the top-of-screen position that the finger would reach if it kept moving in its current direction. As a measure of the online decision, we used either the actual horizontal position of the finger, denoted by x, or the implied endpoint. The x coordinates and the iEP are measured on a scale ranging from −1 (left end of screen) to +1 (right end), with x = 0 being at the middle of the screen. The y coordinates use the same scale (1 unit = 9.84 cm).

2.8. Data analysis

We start by considering only the participants' final responses (left or right). Performance scores are given as the proportion of *ideal responses*, i.e. responses where the participant chose the side that agrees with $sign(\sum WOE)$. Note that, because of the probabilistic nature of the task, the rewarded button was not always the ideal response. However, choosing the ideal response maximized the likelihood to obtain a reward.

Logistic regressions were computed on the proportion of rightward responses, with one predictor corresponding to each of the symbols in the sequence.

The analysis of trajectories followed the method introduced in our earlier publications (Dotan & Dehaene, 2013). The dependent variable was the implied endpoint (*iEP*) and the predictors indicated the different symbols in the sequence (details about specific predictors appear in the text below). One regression was run for each participant and time point, in 30 ms intervals. To determine whether a predictor had a significant group-level effect at a specific time point, we compared, for each time point and each predictor, the per-participant regression coefficients to zero using a *t*-test.

3. Results

3.1. Day 1. Acquisition of the semantic elements of the mini-language

3.1.1. Analysis of the responses

During Day 1, the participants saw only content symbols. We first analyzed the participants' final decision, i.e. whether they ultimately hit the left or right response button. When seeing length-1 sequences, after only a few trials, participants learned to associate each symbol to the corresponding response. The percentage of ideal responses increased from 83% in the first half of Part 1 to 90% in the second half (paired $t(26) = 5.20$, $p < 0.001$).

For length-2 and length-4 sequences, participants also quickly learned to combine the content symbols. They favored the right response button when the total *WOE* (ΣWOE) was positive and the left response button when ΣWOE was negative (Fig. 2a). The proportion of ideal responses was higher in trials with larger ΣWOE, but participants still performed better than chance even on trials in which ΣWOE was close to zero. We fitted the proportion of 'right' responses as a sigmoidal function of ΣWOE:

$$\log_{10}\left(\frac{P^{right}}{1-P^{right}}\right) = const + b \times \sum WOE$$

In this regression, the intercept reflects a general leftward or rightward bias, and the slope measures the participants' capacity to accurately base their decision on ΣWOE. The higher the slope, the closer to optimality: a participant with infinite slope would base his decision on $sign(\sum WOE)$. As the slope of the sigmoidal fit was significantly greater than 0, we conclude that they chose the ideal response determined by $sign(\sum WOE)$ more often for larger $|\Sigma WOE|$ (for length-2 sequences, mean b = 0.67, $t(26) = 9.93$, $p < 0.001$; for length-4 sequence, mean b = 0.45, $t(23) = 12.0$, $p < 0.001$). Furthermore, participants performed better for length-2 than for length-4 sequences, as the slope was higher for length-2 than for length-4 sequences (paired $t(23) = 4.7$, $p < 0.001$). There was no left or right bias – the intercepts were not significantly different from zero, neither for length-2 nor for length-4 sequences (all $p > 0.1$).

To determine whether participants assigned a larger subjective weight to the most reliable symbols (e.g. a larger weight to $>\ >\ >$ than to $>$), we performed a linear regression on the log odds of "right"

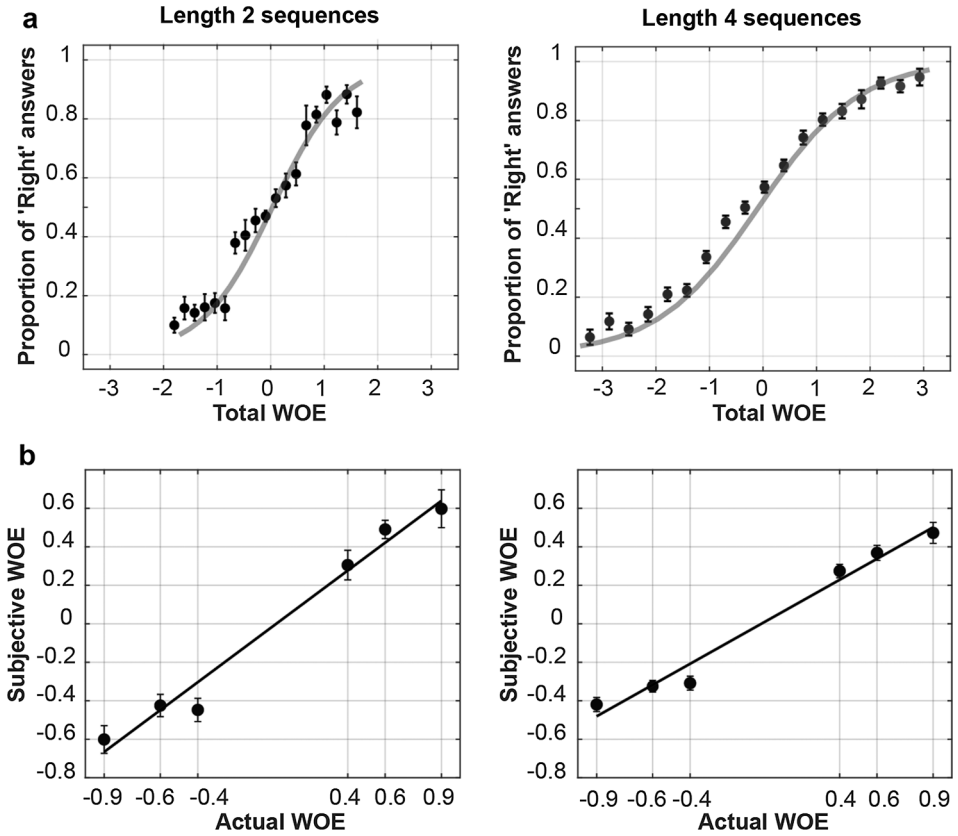

Fig. 2. Evidence that the participants' final decision in Day 1 resulted from a summation of the symbols' WOEs. (a) Proportion of responses to the 'right' response button as a function of the total *WOE*. A sigmoidal fit was obtained for the average regression coefficients over participants. Error bars indicate the standard error on the estimation of the mean (SEM). (b) Average subjective *WOE* as a function of the true *WOE*, as determined by regression. The positive regression slope indicates that participants learned the WOEs and managed to combine them.

responses. Unlike the previous analysis, in which Σ*WOE* was the only predictor, here there was one predictor for each of the 6 semantic symbols. These predictors were defined as 1 when the corresponding symbol was present in the sequence, and as 0 otherwise. The regression coefficients thus reflect the contribution of each symbol to the participant's final decision. Below, we refer to these regression coefficients as the "subjective *WOE*" of the symbol. To determine whether the participants based their decision on the *WOE* of each symbol, we examined how well their subjective *WOE*s were correlated with the true *WOE*. This was done with linear regression on the per-participant subjective *WOE* as a function of the actual *WOE* (Fig. 2b). In this regression, the intercept did not significantly differ from 0 ($p > 0.3$ both in length-2 and in length-4 sequences), meaning that there was no general leftward or rightward bias. The slope was significantly higher than 0 (length-2 sequences: mean b = 0.76, t(26) = 9.13, $p < 0.001$; length-4 sequences: mean b = 0.53, t(23) = 12.2, $p < 0.001$). The order of subjective *WOE* agreed with that of the actual *WOE*s, indicating that participants were able to attribute a higher weight to the symbols that were more strongly associated with a reward. Furthermore, the participants' behavior suggests that they estimated the *WOE* of the symbols better for length-2 than for length-4 sequences: the slope was significantly closer to 1 in the length-2 sequences than in the length-4 sequences (paired t(23) = 3.50, $p = 0.002$).

3.1.2. Analysis of the finger trajectories

Examination of the raw trajectories suggested that they contained rich information on the processing of the sequences of symbols (Fig. 3a) as they seemed to strongly depend on the presented sequences of symbols. This effect can be seen in Fig. 3b, which shows the finger horizontal position (x coordinate) as a function of time, separately for each unique sequence of *sign*(*WOE*) (e.g. the sequences > , > > and > > > , > were averaged together into a single line in Fig. 3b, and separately from the sequence > , < <).

We looked for the time point where each symbol started affecting the finger movement. First we ran, for each position in the sequence (1st, 2nd, 3rd, or 4th symbol), a two-way repeated measures ANOVA with the x coordinate as the dependent variable and with two within-participant factors: the symbol's *sign*(*WOE*), and the specific sequence of *sign*(*WOE*) of the preceding symbols. One ANOVA was run for each time point, in 10 ms intervals. The onset time of a symbol's effect on the finger movement was defined as the first time-point in which this ANOVA showed a significant main effect of *sign*(*WOE*) ($p < 0.05$), which remained significant for at least 50 ms. This analysis revealed that, for length-4 sequences, each successive symbol started affecting the finger movement about 600 ms after its onset (see Fig. 3b). The average interval between the effects of consecutive symbols was ~330 ms, closely corresponding to the actual SOA of 300 ms. These results indicate that the symbols were considered one by one as they appeared, and affected the finger movement even before the full sequence was presented. The finger movements were thus sensitive to intermediate accumulation of evidence.

We next examined how the effect of each symbol developed during a trial and affected the finger trajectory. For each time point and participant, we ran a linear regression on the implied endpoint (*iEP*). There was one predictor for each position – the *WOE* of the symbol at that position. This analysis (Fig. 3c) showed that the effects of each symbol built up gradually over the course of the trial. For length-4 sequences, each symbol had a significant effect on the trajectory ($p < 0.05$ when comparing the per-participant regression coefficients to 0) starting from about 500 ms after the symbol's appearance. The average interval

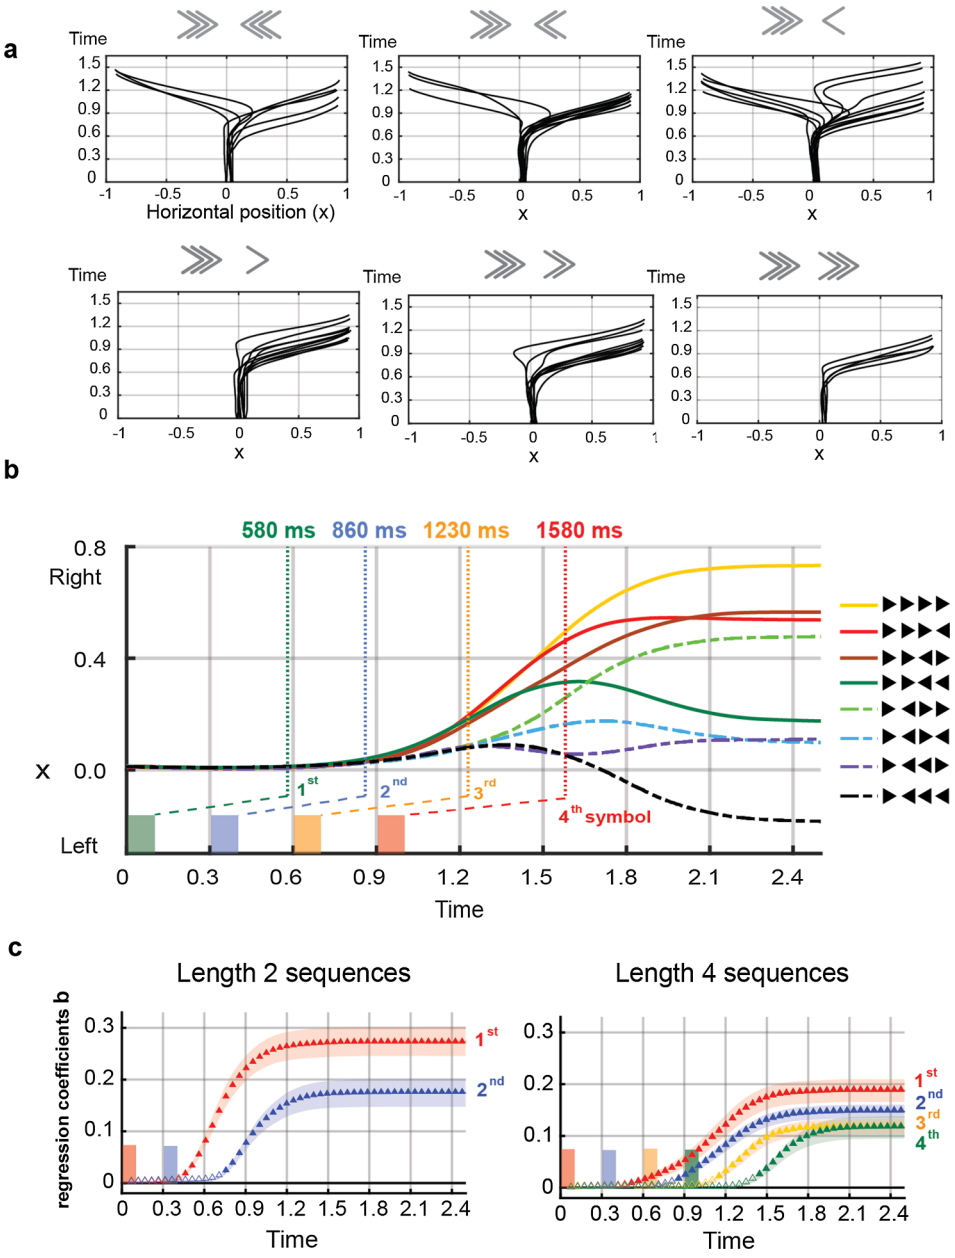

Fig. 3. Finger trajectories in Day 1 reveal a sequential accumulation of evidence. (a) Sample single-trial trajectories for one participant in length-2 sequences. The stimulus is indicated on the top of each panel. The raw trajectories strongly depend on the *WOEs* of both symbols. (b) Mean x coordinate per time point. Each line shows the average trajectory for trials with the same sequence of sign (*WOE*)s. The colored rectangles at the bottom show each symbol's onset time and duration. The colored vertical lines denote the times at which each symbol started having a significant effect on the trajectory. (c) The sequential accumulation of evidence is also captured by regression results. The implied endpoint (*iEP* cf. Fig. 1c) were regressed against the WOE of each position. The regression coefficients were averaged over participants and plotted as a function of time (filled markers indicate values significantly higher than 0). The plots show how, depending on their position in the sequence, symbols successively contribute to the decision. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

between the effects of consecutive symbols was close to the actual SOA (300 ms). These results are consistent with the previous analysis of the symbols' effect onset times and indicate that evidence was accumulated incrementally.

Remarkably, the asymptote values of the regression coefficients decreased with symbol position (Fig. 3c), i.e. earlier symbols affected the finger movement more than symbols that appeared later in the sequence. The contribution of the symbols of the sequence to the intermediate decision, captured by these regression coefficients, probably reflects the contribution of several mechanisms. First of all, this analysis was run on all the length-2 trials of Day 1, while participants were still learning the *WOE*s. The down-weighting of the late symbols in the sequence may thus be partly due to the difficulty of memorization and calculation with symbols that have an uncertain and noisy meaning. However, this explanation may only partially account for this observation, as a similar effect was observed in an experiment where the symbols were not arbitrary shapes but arrows clearly pointing left or right (Dotan et al., 2018). Another explanation comes from the fact that when the *WOE* of the first symbol is large enough, a decision can be reached without considering the next symbols. Once enough evidence is accumulated, an ideal observer may ignore the later symbols, which are not going to change the outcome. For length-2 sequences, this situation occurred 33% of the trials, so for an ideal observer knowing the *WOE*s of the symbols, it should result in the second symbol's regression

coefficient being ⅔ the size of the first.

To determine if this strategy could explain the regression coefficients profiles, we ran the previous analysis on two subsets of trials: those for which enough evidence was provided by the first symbol, and those for which it was not. Fig. S2 shows that the regression coefficient associated with the second symbol was higher when the second symbol was necessary to determine the ideal response (two-tailed paired *t*-test on second position regression coefficients depending on first-position weight: $p < 0.05$ from 0.91 s on). Nonetheless, even when enough evidence was provided by the first symbol, the last symbol still had a sizable influence on the decision.

A logistic regression on the final decision confirmed these results: the regression coefficient of the second position significantly decreased as the evidence carried by the first symbol increased (paired *t*-test between the distributions of regression coefficients when $|WOE|$ at position 1 is 0.4 and when $|WOE|$ at position 1 is 0.9: $t(26) = -3.5$, $p = 0.0015$; when $|WOE|$ at position 1 is 0.6 and when $|WOE|$ at position 1 is 0.9: $t(26) = -3.7$, $p = 0.001$). The converse analysis did not show any significant modulation of the first position regression coefficient as a function of the second symbol's weight ($p > 0.5$), confirming the importance of the temporal constrains of this task in the decision process. To sum up, the participants' behavior was logical (they assigned more weight to earlier symbols), but the degree of this underweighting was smaller than that expected from an optimal observer. It is possible that participants did not manage to adopt an optimal strategy due to the high presentation speed, which may have been too fast to respond as soon as evidence sufficed to determine the ideal response (conscious strategical stopping, de Lange, van Gaal, Lamme, & Dehaene, 2011). Moreover, as this analysis captures the learning stages of the semantic symbols, subjects may still have a noisy representation of the WOE of the different symbols. In this context, the best strategy to compensate for the uncertainty of the task would be to get a maximum of information by waiting for more symbols, which is apparently what participants do.

Finally, the results were remarkably similar to the ones obtained in Dotan et al. (2018) where the stimuli were sequences of arrows pointing left or right and the participants had to point according to the majority of arrows. Even in Dotan et al.'s task, where the *WOE* of each symbol in the sequence ($+\infty$ or $-\infty$, where ∞ denotes infinity) was clearly indicated by the arrow, the earlier symbols in the sequence were over-weighted relatively to later symbols.

3.2. Day 2: Acquisition of the syntactic operator

In the post-experiment questionnaire, 15 out of 25 participants reported the correct interpretation of the operator ⋆, i.e. that it changes the sign of the *WOE* of the subsequent symbol, thus inverting its meaning. From now on, the syntactic operator will equivalently be called inversion operator. This group of participants will be denoted as $G^{\star+}$. In the complementary group, denoted as $G^{\star-}$, 8 participants reported that the ⋆ symbol was different from the other symbols but said that the difference was that ⋆ was not associated with any preferred direction of response (i.e. that it had $WOE = 0$). The remaining two $G^{\star-}$ participants said that ⋆ favored a certain direction of response with a small weight, like the 6 other semantic symbols.

3.2.1. Analysis of the responses

If the participant's behavior was consistent with the questionnaire answers, the $G^{\star-}$ group should attribute $WOE = 0$ to the ⋆ symbol, whereas the $G^{\star+}$ group should treat ⋆ as a inversion operator. To examine whether this was the case, we considered '*critical trials*'. These are trials where correctly interpreting ⋆ as changing the sign of the subsequent WOE would lead to one response (the ideal response), whereas interpreting ⋆ as having $WOE = 0$ would lead to the opposite response. In the $G^{\star+}$ group, the mean accuracy in the '*critical trials*' (70.6%) was significantly better than chance (two-tailed *t*-test to chance level, $t(14) = 6.4$, $p < 0.001$), indicating that these participants correctly interpreted ⋆ as changing the sign of the *WOE* of the subsequence symbol. Conversely, the $G^{\star-}$ group performed on average at chance level for such 2-symbol sequences (45%). Thus, consistently with their explicit report, they did not treat ⋆ as an inversion operator. However, because their performance was not significantly lower than the chance level of 50% (mean accuracy: 45.0%, one-tailed *t*-test to chance level, $t(10) = -1.23$, $p > 0.1$), we cannot conclude that they adopted the $WOE = 0$ interpretation consistently across the block. For length-4 trials, the $G^{\star-}$ participants answered below chance, in agreement with the $WOE = 0$ hypothesis (mean accuracy = 37.4%, one-tailed *t*-test to chance level of 50%, $t(9) = -5.1$, $p < 0.001$). In contrast, the $G^{\star+}$ participants performed at chance level (mean accuracy 50.4%, *t*-test to chance level of 50%: $t(12) = 0.1$, $p > 0.1$) – i.e. they did not fully treat ⋆ as an inversion operator. As we will elaborate in the Discussion, the reason for this may be that long sequences were too hard for them to apply the inversion operation.

To understand in more details how the participants interpreted the syntactic operator, we investigated the contribution of each symbol to the decision. This analysis was first run for length-2 sequences: with logistic regression, the log odds of a rightward response were regressed against 4 predictors: the presence of ⋆ (coded as 0 or 1), the *WOE* of the semantic symbol at position 1 (▲), the *WOE* of the symbol at position 2 when preceded by ⋆ (denoted ⋆▲) and the *WOE* of the symbol at position 2 when preceded by a content symbol (denoted ▲▲). For participants who understand the symbols fully – the inversion operator, the meaning of the content symbols, and how to combine them – we expect regression coefficients b[▲▲] > 0 (i.e. processing the standard *WOE* when the inversion operator did not appear), b[⋆▲] < 0 (i.e. applying the inversion implied by the operator), and b[⋆] = 0 (understanding that the syntactic operator does not have any weight of its own). This was the case in the $G^{\star+}$ group, as shown in Fig. 4a. First, b[⋆] did not significantly differ from 0 ($p > 0.3$). Second, b[▲▲] was higher than zero (mean b = 1.29, $t(14) = 7.2$, $p < 0.001$), indicating that the change of sign of the *WOE* was specifically triggered by the ⋆ operator, and in its absence the participants accumulated the *WOE* normally. Third, crucially, b[⋆▲] was lower than 0 (mean b = −0.6, $t(14) = -6.11$, $p < 0.001$), indicating that these participants changed the sign of the *WOE* of the symbol that appeared after ⋆. We recomputed these regression coefficients on the second half of the block, once the syntactic operator had been learned (*cf* learning curves in Fig. S3). b[⋆] still did not significantly differ from 0 ($p > 0.9$), b[▲▲] was higher than zero (mean b = 1.78, $t(14) = 4.29$, $p < 0.001$) and b[⋆▲] was even more negative (mean b = −1.17, $t(14) = -6.34$, $p < 0.001$), confirming that subjects processed the syntactic operator better in the second half of the block than in the first.

In the $G^{\star-}$ group, interpreting ⋆ as having $WOE = 0$ should lead to b[⋆] = 0, which was indeed the case (mean b[⋆] = −0.25, $t(10) = -1.0$, $p > 0.3$; Fig. 4b). Subjective WOE = 0 should also lead to b[⋆▲] > 0, because ⋆ would simply be ignored. Fig. 4b shows that this was not the case: b[⋆▲] did not significantly differ from 0 (mean b[⋆▲] = 0.16, $t(10) = 1.3$, $p > 0.1$). This pattern suggests that the participants in this group may have responded randomly for sequences with ⋆ rather than attributed $WOE = 0$ to the ⋆ symbol.

A similar analysis was run on length-4 trials. In the $G^{\star+}$ group (Fig. 4c), contrary to the predictions, b[⋆▲] was not significantly negative (mean b[⋆▲] = 0.036, $t(12) = 0.44$, $p = 0.7$). Still, the ⋆ operator significantly decreased the weight of the subsequent symbol (b[⋆▲] < b[▲▲], $t(12) = 2.84$, $p = 0.02$). Thus, although the results were not as clear as for the 2-symbol sequences, even here the $G^{\star+}$ group considered ⋆ as affecting the subsequent content symbol in a direction opposite to the content symbol's *WOE*. In the $G^{\star-}$ group (Fig. 4d), no significant change of sign of the WOE of the next symbol could be measured: b[⋆▲] was significantly positive (mean = 0.37, $t(9) = 6.8$, $p < 0.001$). Moreover, ⋆ did not reduce the weight of the subsequent symbol: b[⋆▲] was not significantly lower than b[▲▲] (t

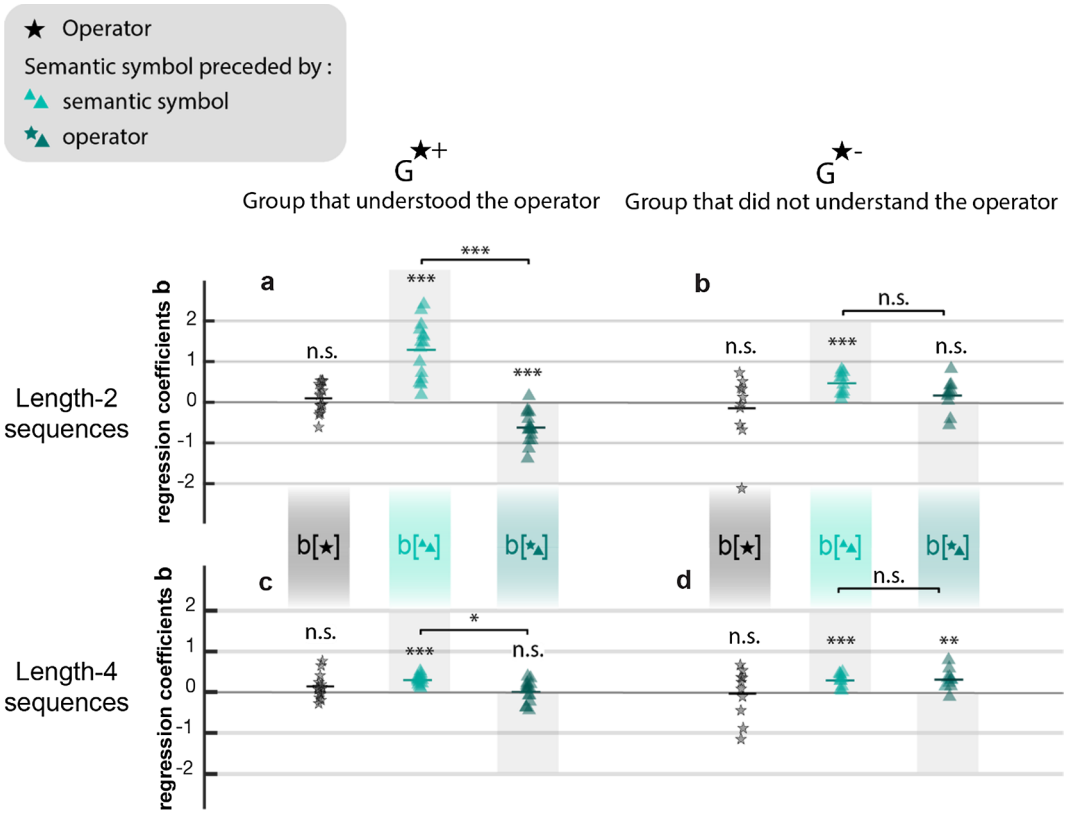

Fig. 4. Understanding of the ★ operator, which changes the sign of the WOE of the subsequent symbol. Two groups of participants were distinguished according to their explicit reports, at the end of the experiment, of understanding ($G^{\star+}$) or failing to understand ($G^{\star-}$) the ★ operator. Each panel shows the weights of a logistic regression on the final decisions during Day 2 for the $G^{\star+}$ and $G^{\star-}$ groups. The regression coefficients presented here are: the presence or absence of ★ (coded as 0 or 1, denoted ★), the *WOE* of the symbol at position 2 when preceded by a content symbol (denoted ▲), and the *WOE* of the symbol at position 2 when preceded by ★ (denoted ★▲). Horizontal segments represent the mean value. Shaded areas represent the range where the regression coefficients should fall for an observer that understands the meaning of the syntactic operator. a. For $G^{\star+}$ participants in length-2 sequences, the operator ★ was not assigned any WOE of its own, but it significantly inverted sign(WOE) of the subsequent symbol ★▲. b. For $G^{\star-}$ participants in length-2 sequences, ★ did not invert sign (WOE) of the subsequent symbol, and did not even reduce the WOE. c. For $G^{\star+}$ participants in the length-4 sequences, ★ decreased significantly the WOE of the subsequent symbol ★▲, yet without inverting its sign. d. For $G^{\star-}$ participants in length-4 sequences, ★ did not change the WOE of the following symbol.

(9) = 0.58, p = 0.6).

3.2.2. Analysis of the finger trajectories

3.2.2.1. Length-2 sequences, $G^{\star+}$ participants. We next analyzed the finger trajectories, starting with the length-2 sequences. We ran linear regressions on the implied endpoints (iEP) in each time point with 4 predictors: the presence of ★ (coded as 0 or 1), the *WOE* of the symbol at position 1 (▲), the *WOE* of the symbol at position 2 when preceded by ★ (★▲), and the *WOE* of the symbol at position 2 when preceded by a content symbol (▲▲). This was done separately for the $G^{\star+}$ and $G^{\star-}$ groups.

The results for length-2 sequences are presented in the top part of Fig. 5. For the subgroup $G^{\star+}$, b[★] was not significantly different from zero in any time point, confirming our earlier conclusion that the participants did not treat ★as a semantic symbol having its own *WOE*. Moreover, consistent with the results of the logistic regression, the regression coefficient associated with ★▲ had a significantly negative value starting from ~400 ms after the appearance of ▲. We note that the onset of b[★▲]'s significant effect was ~150 ms after the onset of b[▲▲]. This 150 ms delay could be a rough estimation of the processing time of the ★ operator. To get a more precise estimation of the operator's processing time, we fitted the regression coefficients to a common exponential profile $A(1 - \exp(-[t - t_0]/\tau))$ for each participant and estimated A, t_0, τ for b[★▲](t) and b[▲▲](t). The parameter t_0, which represents the onset time of a symbol's effect on the finger movement, was significantly higher for b[★▲] than for b[▲▲]. This delay, which is a measure of the processing time of the ★ operator, was estimated as 140 ± 50 ms. No difference was found between b[★▲] and b[▲▲] for the parameter that represents the steepness of the slope (τ = 250 ± 40 ms for both b[★▲] and b[▲▲]).

Initially, the ★operator was novel to the participants. It could be argued that novelty, rather than syntactic function, was at the origin of the observed delay. To provide insight into the learning dynamics of the $G^{\star+}$ group, we analyzed separately the trials belonging to the first and the second half of the length-2 trials of Day 2, the ★ operator being presented in half of the trials of this block (Fig. S4). We observed an increase in the ability to process the syntactic operator, as b[★▲] was much larger in amplitude in the second half of the block ($p < 0.05$ from 800 ms on, uncorrected paired t-test). Moreover, novelty could not account for the ★−induced delay because it was also observed in the second half of the block (the delay is then equal to 120 ± 40 ms), when $G^{\star+}$ participants had already learned the meaning of the ★ operator (*cf* learning curves in Fig. S3). These results suggest that it took the $G^{\star+}$ participants about 120 ms to carry out the inversion indicated by the ★ operator.

3.2.2.2. Length-2 sequences, $G^{\star-}$ participants. The same analysis was ran for the $G^{\star-}$ group and showed that b[★] was not significantly different from zero in any time point, meaning that these participants did not assign the operator any *WOE* of its own. Furthermore, there was no significant difference between b[★▲] and b[▲▲] in any time point (in one-tailed t-test, $p > 0.1$ in all time points), i.e. ★ did not reduce

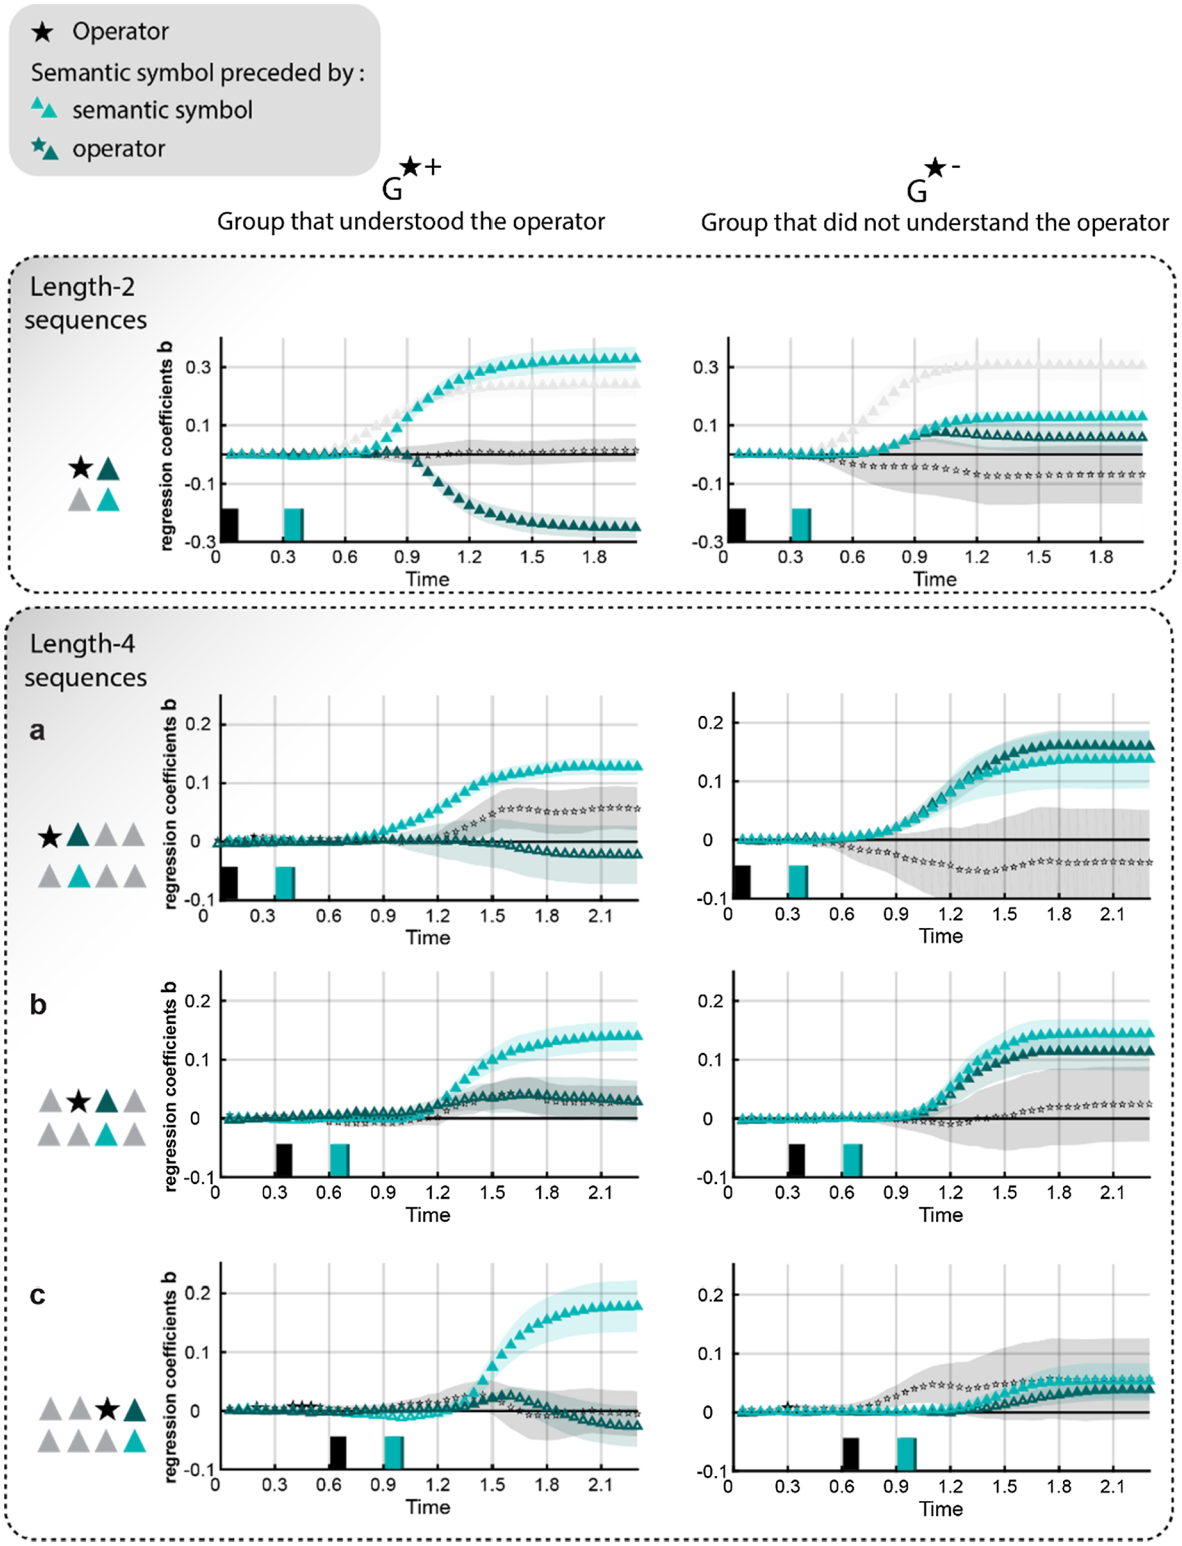

Fig. 5. Finger trajectories reveal how participants processed the syntactic operator. We ran linear regressions on the implied endpoint (*iEP* cf. Fig. 1c) on length-2 and length-4 sequences for the $G^{\star+}$ and $G^{\star-}$ groups (respectively left and right columns). The *iEP* was regressed against the presence or absence of the syntactic operator ⋆ and against the *WOE* of the content symbols at each position in presence or absence of the operator. The shaded areas represent the inter-subject SEM, and filled markers indicate regressors significantly different from 0. For length-2 sequences (top), the regression coefficient b[⋆▲] changed its sign for the $G^{\star+}$ group but not for the $G^{\star-}$ group. Namely, consistently with their report, $G^{\star+}$ participants processed the ⋆ operator as an inversion operation whereas $G^{\star-}$ did not. Fitting with a sigmoid the per-participant regression coefficients, we estimated the processing time of the ⋆-induced inversion operation as 140 ± 50 ms. For length-4 sequences, (bottom), in the $G^{\star+}$ group, but not in the $G^{\star-}$ group, the operator significantly decreased the weight of the subsequent semantic symbol.

the weight of the subsequent symbol. This confirms the results obtained in the logistic regressions: $G^{\star-}$ participants did not process the ⋆ operator as an inversion operator, neither did they consistently assign it a left or right *WOE*.

3.2.2.3. Length-4 sequences. The analysis was extended to length-4 sequences. The dependent variable was the implied endpoint and the 8 predictors were the presence or absence of ⋆ (coded as 1 and 0) and the *WOE* of the content symbols at each position in presence or absence of the operator (Fig. 5 bottom part). For participants in $G^{\star+}$, b[⋆] was not significantly different from zero for any position of ⋆. Similarly, b [⋆▲] was not significantly different from 0 in any time point, except a short time window when ⋆ was in third position. Nonetheless, the operator ⋆ significantly decreased the effect of the subsequent semantic symbol: b[⋆▲] was significantly lower than b[▲▲] ($p < 0.05$ in per time-point one-tailed t-tests) for all positions of ⋆, starting from 1.35 s, 1.30 s and 1.40 s for ⋆ in position 1, 2 and 3, respectively (Fig. 5a–c).

For the participants in the $G^{\star-}$ subgroup, b[⋆] was not significantly different from zero in any time point, for any position of ⋆. Moreover, ⋆ did not significantly reduce the weight of the subsequent symbol: b [⋆▲] and b[▲▲] did not significantly differ in any time point (one-tailed paired *t*-test, $p > 0.05$).

The syntax of our mini-language was such that the ⋆ operator applied only to the semantic symbol that immediately followed it. To determine whether the $G^{\star+}$ participants inappropriately interpreted ⋆ as affecting the other symbols too, we examined their regression coefficients for the content symbols that were at a distance of 2 or 3 from the ⋆ operator (Fig. S1). We could not observe any significant decrease in the regression coefficients when the operator was present (all $p > 0.05$), confirming that $G^{\star+}$ participants understood that ⋆ acted locally, i.e., only on the subsequent symbol.

3.2.3. Generalization of the ⋆ operator

Two content symbols, < < and > >, were preceded by the ⋆ operator only in trials in which feedback was not provided (in Day 2 length-4 trials). This aspect of our design allowed examining whether participants who learned the meaning of ⋆ generalized it to previously-unseen symbol combinations. Such a finding would rule out the possibility that participants learned specific combinations of ⋆ with certain symbols, and would unequivocally show that the participants indeed learned the syntactic meaning of ⋆.

Thus, we now restricted our analysis to trials in which ⋆ was followed by < < or > >. For the $G^{\star+}$ group (2 out of 15 participants did not perform this block), we observed a significant difference b [⋆▲] < b[▲▲] (paired t(12) = 2.45, $p = 0.03$, mean b[⋆▲] = 0.04, mean b[▲▲] = 0.38), i.e. ⋆ decreased the weight of the subsequent symbol. Furthermore, when examining the regression coefficient normalized by the WOE, no significant difference was found between < < and > > and the other symbols (paired t(12) = 1.4, *t*-test $p = 0.2$). Namely, the $G^{\star+}$ participants understood the true syntactic meaning of ⋆ and applied it to previously-unrewarded symbol combinations exactly as they did with the explicitly-rewarded combinations.

For $G^{\star-}$ participants, ⋆ did not significantly decrease the weight of the subsequent symbol (mean b[⋆▲] = 0.44, mean b[▲▲] = 0.41, paired t(9) = 0.45, $p > 0.5$). No significant difference was found between the regression coefficients for < < and > > and the other symbols (paired t(9) = 0.1, $p > 0.5$).

3.2.4. The origins of the differences between the $G^{\star+}$ and $G^{\star-}$ groups

To understand why $G^{\star+}$ and $G^{\star-}$ groups differed in their report and processing of ⋆, we first compared their performance when processing sequences of semantic symbols during Day 1. An inaccurate representation of the WOE of the semantic symbols may make it difficult to learn the ⋆ operator, because participants learn its meaning in conjunction with semantic symbols. No significant difference between the two groups was found in movement time ($p > 0.1$) and performance for length-4 sequences of Day 1 (t(23) = 1.5, $p > 0.1$). However, the $G^{\star-}$ group had significantly fewer ideal responses than the $G^{\star+}$ group for length-1 trials of Day 1 (t(23) = 2.2, one-tailed $p = 0.02$) and length-2 trials of Day 1 (t(23) = 2.0, one-tailed $p = 0.03$). To understand whether the lower accuracy for $G^{\star-}$ group was due to a difficulty in integrating the semantic symbols, we regressed the *iEP* of the length-2 sequences in Day 1 against the weights associated with each position, separately for each group of participants. The $G^{\star-}$ group assigned a much smaller weight to the second symbol than the $G^{\star+}$ group (Fig. S5), suggesting a difficulty in integrating the two symbols. This difficulty may be responsible for their failure to learn the syntactic operator.

This conclusion is further supported by an additional result: we examined whether a participant's ability to integrate the WOE of the semantic symbols correlated with their ability to process the syntactic symbol as an inversion operator. The first ability was quantified as the regression coefficient of the total WOE in a logistic regression on the decision for length-2 sequences of Day 1. The second was quantified as the regression coefficient of a semantic symbol preceded by ⋆ in a logistic regression on the decision for length-2 sequences of Day 2. These two measures correlated ($r^2 = 0.27$, $p < 0.01$), thereby showing a link between WOE-integration and the processing of ⋆.

Additional differences between $G^{\star+}$ and $G^{\star-}$ groups were revealed by the evolution of accuracy while learning the meaning of ⋆. Here, accuracy was defined as the average ideal responses over 30 consecutive '*critical trials*' of the length-2 sequences of Day 2. We remind that the '*critical trials*' are the trials where interpreting * as a content symbol of WOE = 0 would lead to the non-ideal response ⋆. In the first trials, the accuracy of the $G^{\star+}$ group did not differ from chance (mean: 47%, two tailed $p > 0.1$), whereas it was significantly below chance for the $G^{\star-}$ group (mean : 40% , t(9) = −3.0, one-tailed $p < 0.01$). Namely, the $G^{\star-}$ group did not answer randomly, but initially considered ⋆ as a semantic symbol and consistently assigned it a null WOE. For both groups, accuracy increased from the 30 first trials to the 30 last trials ($G^{\star+}$: paired t(14) = 7.0, one-tailed $p < 0.001$, $G^{\star-}$: paired t (9) = −3.5, one-tailed $p < 0.01$), but only the $G^{\star+}$ group ended the block with accuracy above chance level ($G^{\star+}$ mean : 81% , t(14) = 7.9, $p < 0.01$; $G^{\star-}$ mean: 54%, $p > 0.1$). These results indicate that $G^{\star+}$ quickly grasped the syntactic function of ⋆, whereas $G^{\star-}$ seem to have initially considered it as a semantic symbol with null WOE. For a full overview of the dynamics of performance, see Fig. S3.

4. Discussion

We used continuous tracking of finger movement to investigate the covert stages of semantic and syntactic processing in an artificial mini-language. During the first part of the experiment, we investigated how human adults learned the meaning of semantic symbols associated with different weights of evidence favoring left or right response buttons, and the rules to combine sequences of such symbols. All participants understood the semantic component of the artificial mini-language: they understood whether each symbol favored a left or a right response, and their decisions attributed higher weights to symbols that were more predictive of a reward. After a short amount of training, all participants based their decisions on the combined weights of the symbols presented in the sequence. Just like the two rhesus macaques studied in Yang and Shadlen (2007), our participants were able to combine the *WOE*s of the symbols and figure out which response button was the most likely to be rewarded. However, humans required only about an hour of training and ~1,000 trials – much faster than the monkeys, who needed several months and ~130,000 training trials.

Yang and Shadlen (2007) showed that the instantaneous firing rate of LIP neurons in monkeys reflected the evidence accumulated so far, indicating a continuous summation process. The present results show that a similar on-line monitoring of evidence accumulation can be achieved non-invasively, solely using finger trajectory monitoring.

Furthermore, the trajectory data provided information about the effects of discrete symbols at various points within the sequence: for instance, participants assigned less weight to symbols presented later in the sequence, replicating the results from monkeys (Fig. S6 of Yang & Shadlen, 2007). Similar to rhesus macaques, adult humans did not behave ideally. The subjective *WOE* that they assigned to the different symbols underestimated their real weights. Additionally, the weights attributed to the positions in the sequence did not reflect the symbol's optimal contribution to the ideal response. The difficulty of accumulating evidence from several symbols with uncertain meanings in a short amount of time is likely to be at the origin of this non-ideal behavior.

In the second part of the experiment, to investigate the acquisition of a syntactic rule, we incorporated a syntactic operator that reversed the weight of the subsequent symbol. Only 15 out of the 26 participants (the $G^{\bigstar+}$ group) reported to have understood the meaning of this operator. Their subjective introspection was confirmed by objective analyses: in the $G^{\bigstar+}$ group only, left/right responses indicated that the ⋆ symbol was indeed interpreted as inverting the meaning of the subsequent symbol. Finger trajectories showed that in this group, the ⋆ operator significantly affected the subsequent symbol, either changing the sign of its WOE (in short sequences), or at least decreasing its amplitude (in longer sequences). As for the $G^{\bigstar-}$ group, the analyses of decision and of trajectories indicated that, on average, they did not process ⋆ as an inversion operator, but rather as a neutral semantic symbol with $WOE = 0$.

Importantly, the $G^{\bigstar+}$ group did not interpret the ⋆ operator as a global invertor of all subsequent symbols, but correctly applied it only to the symbol following ⋆. Moreover, we ruled out the possibility that the $G^{\bigstar+}$ group learned the meaning of ⋆ by memorizing all combinations of ⋆ with each content symbol: this group managed to generalize the inversion induced by the ⋆ operator to content symbols that were not presented after ⋆ during the learning stages, thereby showing that they truly understood the syntactic meaning of ⋆.

Within the $G^{\bigstar+}$ group, the ⋆ symbol had slightly different effects in the length-2 and length-4 sequences. In the short sequences, it inverted the weight of the subsequent symbol, but in the long sequences it only decreased its weight. This could be explained by assuming two type of trials: presumably, in some trials the participants were able to process the ⋆ operator, whereas in other trials they were not. The regression, which averages the two types of trials, would show a reduced b[⋆▲] rather than a strictly negative b[⋆▲]. This difference between short and long sequences is likely to be due to an overall increase in task difficulty as the sequence becomes longer. Even in the long sequences, the participants' performance would have probably improved had we provided them with further training or longer SOA. However, longer SOAs could have allowed for more explicit strategies, which we wanted to avoid.

The trajectory analysis revealed that in the $G^{\bigstar+}$ group, the ⋆ operator not only inverted the meaning of the subsequent symbol, but also delayed the subsequent symbol's effect on finger movement. Novelty of the ⋆ symbol could not account for this delay as it was still present and estimated to 120 ± 40 ms when we restricted the analysis to the second part of the length-2 trials, i.e. when $G^{\bigstar+}$ participants had already learned the meaning of the operator (*cf* Fig. S4). Furthermore, this was a pure delay – we did not observe a transient activation of the positive value of the symbol, followed by activation of its negative value. Thus, participants who understood the meaning of the inversion operator managed to apply it to the next symbol, but it took about additional 120 ms to invert the symbol before adding its value to the ongoing accumulation of evidence. This finding predicts that, should monkeys be able to learn this mini-language, the firing rate of LIP neurons should also show an additional ~120 ms delay before reflecting the value with opposite sign. As currently designed, this task does not allow us to determine to which extent the processing or memory load associated with the syntactic symbol contributed to the measured delay. Further experiments could explore this aspect by varying the SOA between symbols. Indeed, a refractory period associated to the comprehension of the syntactic operator itself should decrease for large SOA while any delay induced by the application of the inversion operation to the next semantic symbol should stay constant.

By adding the non-commutative operator ⋆, we increased the complexity of the artificial language. We could thus distinguish the human participants that were able to represent syntactic functions, the $G^{\bigstar+}$ group, from those who were only able to attach quantities to symbols, the $G^{\bigstar-}$ group, and who attached a null weight to the operator symbol. Similar intra-species variability was observed in Wilson, Smith, et al. (2015), where only some of the human participants showed sensitivity to the longer-distance relationships present in sequences generated from an artificial grammar with mixed complexity. In our case, the difference in learning the syntactic symbol can be partly explained by differences in the acquisition of the semantic part of the language, i.e. the meaning of the semantic symbols and the combination rule (Section 3.2.4): participants who did not learn the semantic symbols well enough also had difficulty in learning the syntactic symbol. In particular, these participants (the $G^{\bigstar-}$ group) seemed to have initially considered the ⋆ operator as a semantic symbol with null WOE. Furthermore, although these participants did not perform above chance even in the end of the length-2 block in Day 2, they still performed better in the end of the block than in its beginning.

Why were these participants unable to understand the syntactic operator? One possibility is that some minimal mastery of the semantic symbols is required before the syntactic operator can be learned. This suggests that, with additional time, even the participants of the group $G^{\bigstar-}$ would have been able to sharpen their knowledge of semantic symbols and then discover the meaning of ⋆. Another possibility is that they failed to formulate a proper hypothesis about its meaning (e.g. they kept trying to assign it a value rather than a syntactic function). This hypothesis predicts that they would have been able to use this symbol if they had been explicitly told its meaning.

One of the goals of this study was to create a paradigm that allows comparing the mechanisms of syntax acquisition between humans and monkeys, given that monkeys have been shown to be able to acquire the semantic part of this artificial mini-language (Yang & Shadlen, 2007). Nevertheless, we cannot exclude the possibility that our adult human participants simply matched the syntactic operator ⋆ to their pre-existing knowledge of the minus sign in arithmetic. Thus, even if an advantage is observed for human compared to non-human primates, such advantage could reflect prior education rather than a human-unique ability to understand the syntactic operator. This criticism could be addressed by teaching the same language to adults with limited access to formal schooling in mathematics (e.g. Mundurucu Indians; Amalric et al., 2017; Dehaene, Izard, Pica, & Spelke, 2006) or to unschooled children. Alternatively, the present study could be replicated with another syntax unfamiliar to most human adults: the Postfix Polish notation, in which operators follow their operands, may be a good candidate.

Contrary to most artificial grammar studies, our mini-language contains both semantic and syntactic elements and may, as such, involve brain networks similar to the ones of natural language processing. However, the function of our syntactic operator is closer to an arithmetical minus sign (multiplication by -1) than to a linguistic operation of negation (reversal of truth value; Grisoni, Miller, & Pulvermüller, 2017; Kaup & Zwaan, 2003; Tettamanti et al., 2008). It would therefore be interesting to determine, using functional MRI, whether the underlying brain circuits involve classical language areas (left-hemispheric superior temporal sulcus and inferior frontal gyrus) or the distinct math-responsive regions (bilateral parietal, inferior temporal and dorsal prefrontal cortices) (Amalric & Dehaene, 2016, 2017).

Moreover, our mini-language is extremely simple compared to natural languages: it contains a single syntactic operator, whose effect is limited to the subsequent symbol, and does not involve long-distance

relations or embeddings of the kinds that exist in natural language. As such, it may be learnable by non-human primates, since recent studies suggest that they can master grammars that involve local sequential regularities extending to neighboring units (Fitch, 2004; Wilson, Smith, et al., 2015). Wilson, Kikuchi et al. (2015) suggest that the macaque brain regions involved in the processing of such rule-based sequences are the counterparts of the ones associated with the initial stages of syntactic processing in humans (Uhrig, Janssen, Dehaene, & Jarraya, 2016). There is, however, a debate as to whether human-specific language areas such as Broca's area can be activated by simple artificial languages such as the local-global paradigm (Wang et al., 2015) or require more complex human-like center-embedded grammars (Bahlmann, Schubotz, & Friederici, 2008). Indeed, TMS and fMRI studies suggest an involvement of Broca's area (and its right-hemispheric homologue) in complex motor tasks involving sets of hierarchical rules (Alamia et al., 2016; Clerget, Andres, & Olivier, 2013; Clerget, Poncin, Fadiga, & Olivier, 2011; Koechlin & Jubault, 2006). Further studies will determine whether the present paradigm suffices to reveal important differences between human and non-human primates, or whether it should be extended to syntactical operators that act at distance and involve embedding.

Acknowledgements

We thank Pedro Pinheiro-Chagas for his help during the data analysis and Anne Kösem for her suggestions on the manuscript.

This work was supported by the Institut National de la Sante et de la Recherche Medicale (INSERM, http://www.inserm.fr), the Commissariat à l'Energie Atomique et aux Energies Alternatives (CEA, http://www.cea.fr), the Collège de France (https://www.college-de-france.fr/site/college/index.htm), the Bettencourt-Schueller Foundation (http://www.fondationbs.org), and a European Research Council (ERC, https://erc.europa.eu/) grant "Neurosyntax" to S.D.

Declarations of interest

None.

Appendix A. Supplementary material

Supplementary data to this article can be found online at https://doi.org/10.1016/j.cognition.2018.11.006.

References

Alamia, A., Solopchuk, O., D'Ausilio, A., Van Bever, V., Fadiga, L., Olivier, E., & Zénon, A. (2016). Disruption of Broca's area alters higher-order chunking processing during perceptual sequence learning. *Journal of Cognitive Neuroscience, 28*(3), 402–417. https://doi.org/10.1162/jocn_a_00911.

Amalric, M., & Dehaene, S. (2016). Origins of the brain networks for advanced mathematics in expert mathematicians. *Proceedings of the National Academy of Sciences, 113*(18), 4909–4917. https://doi.org/10.1073/pnas.1603205113.

Amalric, M., & Dehaene, S. (2017). Cortical circuits for mathematical knowledge: Evidence for a major subdivision within the brain's semantic networks. *Philosophical Transactions of the Royal Society B: Biological Sciences, 373*(1740), 20160515. https://doi.org/10.1098/rstb.2016.0515.

Amalric, M., Wang, L., Pica, P., Figueira, S., Sigman, M., & Dehaene, S. (2017). The language of geometry: Fast comprehension of geometrical primitives and rules in human adults and preschoolers. *PLoS Computational Biology, 13*(1), https://doi.org/10.1371/journal.pcbi.1005273.

Bahlmann, J., Schubotz, R. I., & Friederici, A. D. (2008). Hierarchical artificial grammar processing engages Broca's area. *NeuroImage, 42*(2), 525–534. https://doi.org/10.1016/j.neuroimage.2008.04.249.

Berthier, N. E. (1996). Learning to reach: A mathematical model. *Developmental Psychology, 32*(5), 811–823. https://doi.org/10.1037/0012-1649.32.5.811.

Clerget, E., Andres, M., & Olivier, E. (2013). Deficit in complex sequence processing after a virtual lesion of left BA45. *PLoS ONE, 8*(6), e63722. https://doi.org/10.1371/journal.pone.0063722.

Clerget, E., Poncin, W., Fadiga, L., & Olivier, E. (2011). Role of Broca's area in implicit motor skill learning : Evidence from continuous theta-burst magnetic stimulation (pp. 80–92). MIT Press. Retrieved from https://www.mitpressjournals.org/doi/abs/10.1162/jocn_a_00108.

de Lange, F. P., van Gaal, S., Lamme, V. A. F., & Dehaene, S. (2011). How awareness changes the relative weights of evidence during human decision-making. *PLoS Biology, 9*(11), https://doi.org/10.1371/journal.pbio.1001203.

Dehaene, S., Izard, V., Pica, P., & Spelke, E. (2006). Core knowledge of geometry in an amazonian group. *Science, 311*(5579), 381–384. https://doi.org/10.1126/science.1121739.

Dotan, D., & Dehaene, S. (2013). How do we convert a number into a finger trajectory? *Cognition, 129*(3), 512–529. https://doi.org/10.1016/j.cognition.2013.07.007.

Dotan, D., Meyniel, F., & Dehaene, S. (2018). On-line confidence monitoring during decision making. *Cognition, 171*, 112–121. https://doi.org/10.1016/j.cognition.2017.11.001.

Erb, C. D., Moher, J., Sobel, D. M., & Song, J. H. (2016). Reach tracking reveals dissociable processes underlying cognitive control. *Cognition, 152*, 114–126. https://doi.org/10.1016/j.cognition.2016.03.015.

Fitch, W. T. (2004). Computational constraints on syntactic processing in a nonhuman primate. *Science, 303*(5656), 377–380. https://doi.org/10.1126/science.1089401.

Frank, M. C., & Tenenbaum, J. B. (2011). Three ideal observer models for rule learning in simple languages. *Cognition, 120*(3), 360–371. https://doi.org/10.1016/j.cognition.2010.10.005.

Friederici, A. D., Steinhauer, K., & Pfeifer, E. (2002). Brain signatures of artificial language processing: Evidence challenging the critical period hypothesis. *Proceedings of the National Academy of Sciences, 99*(1), 529–534. https://doi.org/10.1073/pnas.012611199.

Friedman, J., Brown, S., & Finkbeiner, M. (2013). Linking cognitive and reaching trajectories via intermittent movement control. *Journal of Mathematical Psychology, 57*(1–2), 140–151. https://doi.org/10.1016/j.jmp.2013.06.005.

Grisoni, L., Miller, T. M., & Pulvermüller, F. (2017). Neural correlates of semantic prediction and resolution in sentence processing. *The Journal of Neuroscience, 37*(18), 4848–4858. https://doi.org/10.1523/JNEUROSCI.2800-16.2017.

Hauser, M. D. (2002). The faculty of language: What is it, who has it, and how did it evolve? *Science, 298*(5598), 1569–1579. https://doi.org/10.1126/science.298.5598.1569.

Hauser, M. D., Newport, E. L., & Aslin, R. N. (2001). Segmentation of the speech stream in a non-human primate: Statistical learning in cotton-top tamarins. *Cognition, 78*(3), https://doi.org/10.1016/S0010-0277(00)00132-3.

Kaup, B., & Zwaan, R. A. (2003). Effects of negation and situational presence on the accessibility of text information. *Journal of Experimental Psychology: Learning Memory and Cognition, 29*(3), 439–446. https://doi.org/10.1037/0278-7393.29.3.439.

Koechlin, E., & Jubault, T. (2006). Broca's area and the hierarchical organization of human behavior. Elsevier, 50, 963–974. Retrieved from https://www.sciencedirect.com/science/article/pii/S0896627306004053.

Marcus, G. F., Vijayan, S., Bandi Rao, S., Vishton, P. M., Saffran, J., Aslin, R., … Holyoak, K. J. (1999). Rule learning by seven-month-old infants. *Science (New York, N.Y.), 283*(5398), 77–80. https://doi.org/10.1126/science.283.5398.77.

Milne, A. E., Mueller, J. L., Mannel, C., Attaheri, A., Friederici, A. D., & Petkov, C. I. (2016). Evolutionary origins of non-adjacent sequence processing in primate brain potentials. *Scientific Reports, 6*, 36259. https://doi.org/10.1038/srep36259.

Milne, A. E., Petkov, C. I., & Wilson, B. (2018). Auditory and visual sequence learning in humans and monkeys using an artificial grammar learning paradigm. *Neuroscience*. https://doi.org/10.1016/j.neuroscience.2017.06.059.

Moeser, S. D., & Olson, A. J. (1974). The role of reference in children's acquisition of a miniature artificial language. *Journal of Experimental Child Psychology, 17*(2), 204–218. https://doi.org/10.1016/0022-0965(74)90066-6.

Mueller, J. L., Hahne, A., Fujii, Y., & Friederici, A. D. (2005). Native and nonnative speakers' processing of a miniature version of Japanese as revealed by ERPs. *Journal of Cognitive Neuroscience, 17*(8), 1229–1244. https://doi.org/10.1162/0898929055002463.

Newport, E. L., Hauser, M. D., Spaepen, G., & Aslin, R. N. (2004). Learning at a distance II. Statistical learning of non-adjacent dependencies in a non-human primate. *Cognitive Psychology*. https://doi.org/10.1016/j.cogpsych.2003.12.002.

Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin's mistake: Explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences, 31*(02), https://doi.org/10.1017/S0140525X08003543.

Pinheiro-Chagas, P., Dotan, D., Piazza, M., & Dehaene, S. (2017). Finger tracking reveals the covert processing stages of mental arithmetic. *Open Mind: Discoveries in Cognitive Science,* 1–12. https://doi.org/10.1162/OPMI_a_00003.

Ravignani, A., Sonnweber, R. S., Stobbe, N., & Fitch, W. T. (2013). Action at a distance: Dependency sensitivity in a New World primate. *Biology Letters*. https://doi.org/10.1098/rsbl.2013.0852.

Reber, A. S. (1969). Transfer of syntactic structure in synthetic languages. *Journal of Experimental Psychology, 81*(1), 115–119. https://doi.org/10.1037/h0027454.

Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changing your mind : A computational mechanism of vacillation. *Nature, 461*(7261), 263–266. https://doi.org/10.1038/nature08275.

Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition, 70*(1), 27–52. https://doi.org/10.1016/S0010-0277(98)00075-4.

Song, J. H., & Nakayama, K. (2009). Hidden cognitive states revealed in choice reaching tasks. *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2009.04.009.

Sonnweber, R., Ravignani, A., & Fitch, W. T. (2015). Non-adjacent visual dependency learning in chimpanzees. *Animal Cognition*. https://doi.org/10.1007/s10071-015-0840-x.

ten Cate, C., & Okanoya, K. (2012). Revisiting the syntactic abilities of non-human animals: Natural vocalizations and artificial grammar learning. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 367*(1598), 1984–1994. https://doi.org/10.1098/rstb.2012.0055.

Tettamanti, M., Manenti, R., Della Rosa, P. A., Falini, A., Perani, D., Cappa, S. F., & Moro, A. (2008). Negation in the brain: Modulating action representations. *NeuroImage, 43*(2), 358–367. https://doi.org/10.1016/J.NEUROIMAGE.2008.08.004.

Uhrig, L., Janssen, D., Dehaene, S., & Jarraya, B. (2016). Cerebral responses to local and global auditory novelty under general anesthesia. *NeuroImage, 141*, 326–340. https://doi.org/10.1016/j.neuroimage.2016.08.004.

Wang, L., Uhrig, L., Jarraya, B., & Dehaene, S. (2015). Representation of numerical and sequential patterns in macaque and human brains. *Current Biology, 25*(15), 1966–1974. https://doi.org/10.1016/j.cub.2015.06.035.

Wilson, B., Kikuchi, Y., Sun, L., Hunter, D., Dick, F., Smith, K., ... Petkov, C. I. (2015). Auditory sequence processing reveals evolutionarily conserved regions of frontal cortex in macaques and humans. *Nature Communications, 6*, 8901. https://doi.org/10.1038/ncomms9901.

Wilson, B., Marslen-Wilson, W. D., & Petkov, C. I. (2017). Conserved sequence processing in primate frontal cortex. *Trends in Neurosciences*. https://doi.org/10.1016/j.tins.2016.11.004.

Wilson, B., Smith, K., & Petkov, C. I. (2015). Mixed-complexity artificial grammar learning in humans and macaque monkeys: Evaluating learning strategies. *European Journal of Neuroscience, 41*(5), 568–578. https://doi.org/10.1111/ejn.12834.

Yang, T., & Shadlen, M. N. (2007). Probabilistic reasoning by neurons. *Nature, 447*(7148), 1075–1080. https://doi.org/10.1038/nature05852.