

Towards a more flexible language of thought: Bayesian grammar updates after each concept exposure

Pablo Tano,¹ Sergio Romano,^{1,2} Mariano Sigman,^{3,4,5} Alejo Salles,⁶ and Santiago Figueira^{1,2}

¹*Universidad de Buenos Aires, Facultad de Ciencias Exactas y Naturales, Argentina*

²*CONICET-Universidad de Buenos Aires, Instituto de Ciencias de la Computación, Argentina*

³*Laboratorio de Neurociencia, Universidad Torcuato Di Tella, Buenos Aires, Argentina*

⁴*CONICET (Consejo Nacional de Investigaciones Científicas y Técnicas), Argentina*

⁵*Facultad de Lenguas y Educación, Universidad Nebrija, Madrid, Spain*

⁶*CONICET-Universidad de Buenos Aires, Instituto de Cálculo, Argentina*



(Received 27 September 2019; accepted 16 March 2020; published 23 April 2020)

Recent approaches to human concept learning have successfully combined the power of symbolic, infinitely productive rule systems and statistical learning to explain our ability to learn new concepts from just a few examples. The aim of most of these studies is to reveal the underlying language structuring these representations and providing a general substrate for thought. However, describing a model of thought that is fixed once trained is against the extensive literature that shows how experience shapes concept learning. Here, we ask about the plasticity of these symbolic descriptive languages. We perform a concept learning experiment that demonstrates that humans can change very rapidly the repertoire of symbols they use to identify concepts, by compiling expressions that are frequently used into new symbols of the language. The pattern of concept learning times is accurately described by a Bayesian agent that rationally updates the probability of compiling a new expression according to how useful it has been to compress concepts so far. By portraying the language of thought as a flexible system of rules, we also highlight the difficulties to pin it down empirically.

DOI: [10.1103/PhysRevE.101.042128](https://doi.org/10.1103/PhysRevE.101.042128)

I. INTRODUCTION

How can children acquire a vast universe of concepts with seemingly very little exposure? One possible solution to this conundrum, known as the Plato Problem [1,2], builds on the human capacity to describe concepts—and more generally of all elements of thought—through the use of a symbolic and combinatorial mental language [3], referred as *language of thought* (LoT) [4].

Combinatorial languages can describe a vast set of concepts from a small set of primitives. This can be understood in a relatively simple example in the domain of shapes. A combinatorial and symbolic language similar to Logo [5] can combine operations such as “move,” “pen up,” “pen down,” or “rotate” to generate an infinite set of expressions (or programs) which, when evaluated, can convey all sort of shapes.

A language describing concepts (like shapes) also provides a natural notion of their complexity [6]. A concept is simple, relative to that language, when it can be described by a short program. On the contrary, it is complex when all its descriptions require a long sequence of instructions. For example, in the case of the Logo language, a square can simply be instructed as a loop of four displacements followed by rotations of 90 degrees. In this language, the icon of a face will be implemented by a significant lengthier program and hence will be more complex. However, this concept would be simpler when described in a language in which the icon of a face (or the symbols for nose, mouth, etc.) are available as primitives in the language.

In the domain of Boolean concepts, a wide range of logical varieties of concepts was studied in Ref. [7], revealing a surprisingly simple “law”: the subjective difficulty of a Boolean concept for a human learner is directly proportional to the length of the shortest compatible program in the language of propositional logic (i.e., Boolean variables combined with the operators *and*, *or*, and *not*). This result may suggest that human LoT is equipped with rules and symbols similar to those found in propositional logic. Indeed, the correlation between the subjective difficulty of concepts and their complexity has been used as a general vehicle to study human LoT in various domains [8–12]. Although often implicit, the general strategy is to (1) assume a language; (2) find the shortest compatible program for some concepts in that language; (3) compare the length of these programs with the subjective difficulty of the concepts; and finally (4) repeat this process for various languages within a universe of possible candidates and choose the language that gives the best match in (3). As mentioned before, the length of the program depends on the primitives of the language in which this program is written, so different languages make different predictions.

A natural question, however, is whether the primitives of a LoT are universal—both across different individuals and also throughout development—or if instead the semantic repertoire of a language is dynamic and shaped by experience. Indeed, it is likely that our ability to automatically represent Boolean concepts in a succinct manner is not due to an innate efficient propositional language in our mind. Instead, we propose that this ability arises as a byproduct of our brain rapidly learning

efficient representations for the concepts we usually encounter in everyday life. Our research question is: how rapidly can we adapt our learning mechanisms when we encounter a new domain in which our *a priori* representations are no longer efficient? We examine the hypothesis that humans have the ability to rapidly recombine propositions in their LoT, adding new primitives to their language. In other words, that learning leads to a process of compiling routines into functions within the LoT.

In the example of the Logo language one can imagine that if productions which draw squares are very frequent, it would be efficient to devote a new symbol to this production. The new symbol “square” is a hierarchical “second-order” construction of the “first-order” primitives of the language. It has a cost (of increasing the lexicon of the language) but in the new language, drawing a square can be instantiated with a very short program (namely, “square”) and hence uses less memory. Indeed, a higher level language allows us to reach a higher level of abstraction by freeing memory and processing power, thus making more complex thoughts thinkable [13,14].

Most work in the LoT literature, while naturally including a learning mechanism, tends to approach the LoT as a stable system to be unearthed by experimenters, who try different candidate templates and select the one which best fits the data after training [8,15,16]. Still, how different tracks of experience can shape acquisition differently and can constantly change the repertoire of a LoT after each exposure remains to be discovered.

Here, we perform a Boolean concept learning experiment to show that humans can change very rapidly—in the course of an experiment—the repertoire of symbols they use to identify concepts. We also provide a dynamic model that is flexible enough to update its underlying language after each concept exposure.

In our experiment, participants are divided in two groups, in such a way that each group is presented with a different sequence of concepts. One of the two groups is presented with concepts that are succinctly described only if the logical operator “exclusive or” (xor, notated \oplus) is used, which we presume does not form part of the natural repertoire of LoT in this specific domain [8]. However, these concepts can also be described with a sensibly lengthy combination of primitives excluding \oplus . We show how the exposure to this set of concepts “compiles” the \oplus operator in a way that, after exposure, subjective difficulty is described by an extended language in which \oplus has been incorporated to the set of primitives. Furthermore, we show that the subjective difficulty of concepts throughout the task is consistent with that of a Bayesian agent that rationally updates the probability of compiling \oplus according to how useful it has been to compress concepts so far.

II. THE LOGICAL SETTING

We consider two propositional logics, both containing only four propositional variables $\mathbf{Vars} = \{x_1, x_2, x_3, x_4\}$. \mathbf{P} is defined over the signature \wedge, \vee, \neg , and \mathbf{P}^\oplus is defined over the signature \wedge, \vee, \neg and \oplus . As one can see from the grammars defined in Fig. 1, the only difference between \mathbf{P} and \mathbf{P}^\oplus is that the latter has an additional operator \oplus .

START \rightarrow BOOL	For $i = 1, 2, 3, 4$
BOOL \rightarrow (BOOL \wedge BOOL)	ATOM $\rightarrow x_i$
BOOL \rightarrow (BOOL \vee BOOL)	ATOM $\rightarrow \neg x_i$
BOOL \rightarrow ATOM	

FIG. 1. The context-free grammar for language \mathbf{P} . Language \mathbf{P}^\oplus has an extra rule: $\text{BOOL} \rightarrow (\text{BOOL} \oplus \text{BOOL})$.

The semantics of $\wedge, \vee,$ and \neg are standard: conjunction, disjunction, and negation, respectively. We let \oplus denote the exclusive disjunction. As usual, $v \models \varphi$ represents that the formula φ is true for the valuation $v : \mathbf{Vars} \rightarrow \{0, 1\}$ and we denote the semantics of φ by $\llbracket \varphi \rrbracket = \{v : v \models \varphi\}$. A concept \mathcal{C} is a set of valuations $\mathbf{Vars} \rightarrow \{0, 1\}$. The complement of \mathcal{C} is denoted $\bar{\mathcal{C}}$ and is defined as $\bar{\mathcal{C}} = \{0, 1\}^{\mathbf{Vars}} \setminus \mathcal{C}$. Observe that $\#\mathcal{C} + \#\bar{\mathcal{C}} = 16$. We say that a formula φ is compatible with concept \mathcal{C} if $\llbracket \varphi \rrbracket = \mathcal{C}$. We regard logics as languages for describing concepts. Any concept \mathcal{C} has infinitely many descriptions, namely, all formulas φ such that $\llbracket \varphi \rrbracket = \mathcal{C}$.

Example. In Fig. 2 we depict a concept \mathcal{C} (variables are represented by colors) such that $\#\mathcal{C} = 4$. One can see that the formula x_3 is not compatible with \mathcal{C} but $x_1 \wedge x_2$, or $x_1 \wedge x_2 \wedge (x_3 \vee \neg x_3)$, are compatible with \mathcal{C} . $\bar{\mathcal{C}}$ may be described by $\neg x_1 \vee \neg x_2$.

We will often identify concepts with any formula compatible with it, so we will talk of “concept φ ” to refer to “concept $\llbracket \varphi \rrbracket$.” However, it should be noted that a concept is a semantic object that has many descriptions in the logical language.

A lower bound for the complexity of a concept in a given logic corresponds to the shortest description of that concept, that is, its minimum description length (MDL).

The size of a formula φ is denoted $|\varphi|$ and it is defined as the number of operators plus the number of atoms (i.e., possibly negated propositional symbols), that is: $|x_i| = |\neg x_i| = 1$ for $i = 1 \dots 4$ and $|\varphi_1 * \varphi_2| = |\varphi_1| + |\varphi_2| + 1$ for $*$ $\in \{\wedge, \vee, \oplus\}$. For $\mathbf{L} \in \{\mathbf{P}, \mathbf{P}^\oplus\}$ and a concept \mathcal{C} we define the minimum description length of \mathcal{C} with respect to \mathbf{L} as

$$\text{MDL}_{\mathbf{L}}(\mathcal{C}) = \min\{|\varphi| : \varphi \in \mathcal{L}, \llbracket \varphi \rrbracket = \mathcal{C}\}.$$

Since \mathbf{P} is a sublanguage of \mathbf{P}^\oplus , we have $\text{MDL}_{\mathbf{P}^\oplus}(\mathcal{C}) \leq \text{MDL}_{\mathbf{P}}(\mathcal{C})$ for any concept \mathcal{C} .

Example. The concept $\mathcal{C} = \{v : v(x_1) + v(x_2) = 1\}$, which expresses that x_1 is true or x_2 is true but not both can be described in \mathbf{P}^\oplus as $\varphi = x_1 \oplus x_2$, of length 3. In fact, one can check that this is the shortest formula compatible with \mathcal{C} , and so $\text{MDL}_{\mathbf{P}^\oplus}(\mathcal{C}) = 3$. If we now switch to \mathbf{P} , then we




FIG. 2. Example of a concept \mathcal{C} , as shown in the experiment. Variables in $\mathbf{Vars} = \{x_1, x_2, x_3, x_4\}$ correspond to the presence of a color light in the object ($x_1 = \text{green}, x_2 = \text{blue}, x_3 = \text{red}, x_4 = \text{orange}$). Items (valuations) belonging to \mathcal{C} are highlighted with bold border. \mathcal{C} may be described by $x_1 \wedge x_2$. As in a traffic light, each color is fixed to each position.

can no longer describe \mathcal{C} as $x_1 \oplus x_2$, since \oplus is not part of its signature. However, in \mathbf{P} , the concept \mathcal{C} may be described by formula $\psi = (x_1 \wedge \neg x_2) \vee (x_2 \wedge \neg x_1)$, of size 7. Since this formula has minimal size, we have that $\text{MDL}_{\mathbf{P}}(\mathcal{C}) = 7$.

III. EXPERIMENT

Fifty-five participants participated in the experiment over the world wide web using the Amazon Mechanical Turk crowd-sourcing platform. All were U.S. residents over the age of 18 and had more than 95% of past tasks successfully approved by other requesters. Forty-four participants completed the experiment through all the stages and declared not cheating (using pen, screenshots, or a similar method to copy the answers) at the end of the experiment. Only data from these participants were used in the analyses reported below.¹

Participants were divided randomly into a control group ($N = 21$) and a target group ($N = 23$). Both groups were presented with different sequences of six concepts. For each concept, there was a learning phase, a testing phase and a feedback phase. The average time spent in each concept was 167 ± 20 standard error of the mean (s.e.m.) seconds, and the average duration of the task was 21 ± 4 s.e.m. minutes. After moving through the learning, testing, and feedback phase of each of the six concepts, participants were asked if they used a pen or recorded the screen information in any way. They were also told that the answer to this question will not affect their payment, but that it was crucial for the experimenters to know.

During the learning phase, all 16 items were presented in the screen (in random order), and items belonging to the concept were identified with bold boundaries, as shown in Fig. 2. Participants were told that only the items with bold boundaries were “blickets” (or “tufas,” etc.: we used different words for each concept in the sequence), and asked them to try to identify what a blicket was. During the testing phase, the 16 items were shuffled in the screen, and participants were asked to click on items that were blickets. If they made mistakes after submitting their answer, then they were directed to the feedback phase, in which items that were incorrectly classified were indicated with a red cross. After having studied the feedback, participants were redirected to the testing screen, where items were reshuffled. When every item was correctly classified, participants were asked to give a verbal description

of the concept and then continued on to the following concept after a resting period. We characterize the subjective difficulty of each concept as the time the participant spent in learning, testing and feedback phases for that concept (excluding the time spent in the verbal description).

Both groups (target and control), were exposed to six concepts. The second, third, and fourth concepts are *training* concepts, and were different between both groups. The last two concepts are the *test* concepts and were the same for both groups. The first concept was the trivial concept x_i for both groups, which was aimed to get participants started in the task. Importantly, variables (i.e., color lights inside objects in Fig. 2) were randomized for every concept, so paying selective attention to a specific variable across subsequent concepts was not beneficial for learning the concept sequence.

As shown in Table I, we presented the target group with training concepts which are succinctly described when \oplus is part of the language, but necessarily described with lengthier formulas if \oplus is absent; more technically, concepts for which $\text{MDL}_{\mathbf{P}^\oplus}$ is much smaller than $\text{MDL}_{\mathbf{P}}$. We also corroborated that for \mathcal{C}_t^2 , \mathcal{C}_t^3 , \mathcal{C}_t^4 , and \mathcal{C}^5 the number of formulas in \mathbf{P}^\oplus with length strictly smaller than $\text{MDL}_{\mathbf{P}}$ was at least 10 times greater than the number of formulas in \mathbf{P} with length equal to $\text{MDL}_{\mathbf{P}}$.

Participants in the control group, on the other hand, experienced a sequence of concepts that could be easily described using the language given by \mathbf{P} . After these training concepts, both groups were presented with the same pair of test concepts: one which could be only succinctly described in \mathbf{P}^\oplus , and one for which the MDL did not depend on the underlying language \mathbf{P}^\oplus or \mathbf{P} . We compared learning times between the two groups for these last two concepts.

As shown in Table I, training concepts for the target (xor) group were: x_i , $x_i \oplus x_j$, $x_i \oplus x_j \oplus x_k$, and $x_k \oplus x_l$, called \mathcal{C}_t^1 , \mathcal{C}_t^2 , \mathcal{C}_t^3 , and \mathcal{C}_t^4 , respectively. Training concepts for the control group were: x_i , $x_i \vee x_j$, $x_i \vee (x_j \wedge x_k)$, and $x_k \vee x_l$ called \mathcal{C}_c^1 , \mathcal{C}_c^2 , \mathcal{C}_c^3 , and \mathcal{C}_c^4 , respectively. We use the indexes i, j, k, l instead of numbers because variables were randomized in each trial. x_i could stand for x_1, x_2, x_3 or x_4 , that is, for any of the four colors. After these four concepts, both groups were presented with the same test concepts: $x_i \wedge (x_j \oplus x_k)$ and $x_i \wedge (x_j \vee x_k)$, called \mathcal{C}^5 and \mathcal{C}^6 , respectively.

Choosing which concepts to show the target group in order for them to “learn” the \oplus operator is critical in our experiment. Crucially, the learner must have an option between two alternatives that describe the concept: one that is succinct but uses \oplus , or necessarily a much longer one in the absence of \oplus . In other words, these concepts must be compatible with short

¹The learning times of all participants can be found in Ref. [17].

TABLE I. Sequence of concepts presented in the experiment: $\mathcal{C}^1, \mathcal{C}_t^2, \mathcal{C}_t^3, \mathcal{C}_t^4, \mathcal{C}^5, \mathcal{C}^6$ for target group and $\mathcal{C}^1, \mathcal{C}_c^2, \mathcal{C}_c^3, \mathcal{C}_c^4, \mathcal{C}^5, \mathcal{C}^6$ for control group. Each concept \mathcal{C} is represented by a minimal formula φ such that $\llbracket \varphi \rrbracket = \mathcal{C}$.

	Target group		$\text{MDL}_{\mathbf{P}^\oplus}(\mathcal{C})$	$\text{MDL}_{\mathbf{P}}(\mathcal{C})$	Control group		$\text{MDL}_{\mathbf{P}^\oplus}(\mathcal{C})$	$\text{MDL}_{\mathbf{P}}(\mathcal{C})$
Training	\mathcal{C}^1	x_i	1	1			←Idem	
	\mathcal{C}_t^2	$x_i \oplus x_j$	3	7	\mathcal{C}_c^2	$x_i \vee x_j$	3	3
	\mathcal{C}_t^3	$x_i \oplus x_j \oplus x_k$	5	19	\mathcal{C}_c^3	$x_i \vee (x_j \wedge x_k)$	5	5
	\mathcal{C}_t^4	$x_k \oplus x_l$	3	7	\mathcal{C}_c^4	$x_k \vee x_l$	3	3
Test	\mathcal{C}^5	$x_i \wedge (x_j \oplus x_k)$	5	9			←Idem	
	\mathcal{C}^6	$x_i \wedge (x_j \vee x_k)$	5	5			←Idem	

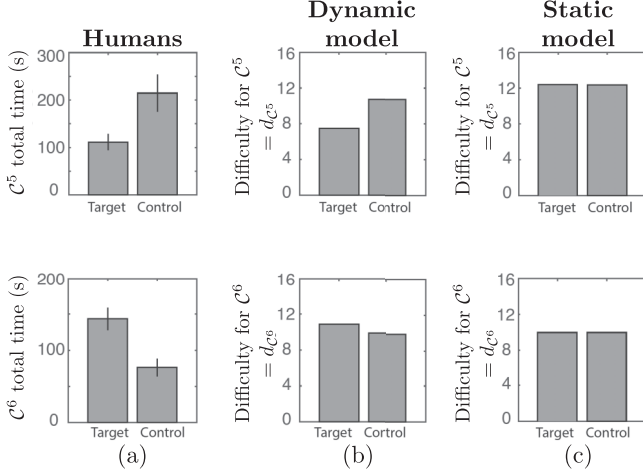


FIG. 3. Concept learning time (a) and difficulty predicted (b, c) for the two test concepts (C^5 and C^6). Error bars are s.e.m. across subjects.

logical formulas if and only if we take P^\oplus as the language of description. To ensure that this was the case, we enumerated, for each concept, all formulas compatible with it and produced by the P and P^\oplus grammars up to length 19. For all training concepts of the target group, the shortest compatible formula without \oplus is much longer than the shortest compatible formula with \oplus . This is shown in Table I.

IV. MODEL-FREE RESULTS

We measure the subjective difficulty of a given concept as the total time needed by the participant to successfully encode the concept, which indicates that they can reliably express which exemplars belong to the concept and which do not.

Participants from the target group spent almost half the time than participants from the control group in C^5 , which could be succinctly described only in P^\oplus (111 ± 16 s.e.m. seconds versus 214 ± 37 s.e.m. seconds, a two-sample t -test reveals $t_{42} = 2.6$, $P < 0.01$), as shown in Fig. 3(a). We also found that the control group learned much faster C^6 (143 ± 14 s.e.m. seconds for the target group versus 76 ± 10 s.e.m. seconds for the control group, $t_{42} = 3.5$, $P < 0.01$). A mixed ANOVA with C^5 - C^6 as within subject factor and target-control groups as between subject factor reveals a strong interaction between group and C^5 - C^6 ($F = 15.3$, $P < 0.001$), indicating that the differences in learning times for C^5 and C^6 were very different between the two groups.

The target group then encoded C^5 more efficiently than the control group. We propose that the control group expected to find in C^5 and C^6 structures that could be easily built in P . The target group, on the other hand, became biased towards the \oplus structure, and they expected to find it in C^5 and C^6 . This caused C^5 to be encoded more rapidly by the target group and C^6 more rapidly by the control group. Assuming that the subjective difficulty of learning a concept is proportional to the complexity of its internal representation, we conclude that after exposure to the training concepts, participants in the target group represented the \oplus more efficiently than the control group, and expected to find this structure in C^5 and C^6 .

V. MODEL

When presented with a concept (e.g., Fig. 2), our model generates logical formulas and evaluate them to true or false for that concept, keeping the formula only if it is true. To generate formulas, the model uses a symbolic language in which each rule (symbols and operators) is associated with a probability of being used. The probability of generating a formula is proportional to the product of the probabilities of the rules required for building it, and therefore it decreases exponentially with its length. Furthermore, if one of the rules has a very low probability of being used, then formulas that require it will also have very low probability.

The *Static model* maintains the rules' probabilities fixed throughout the concept sequence (the six concepts in Table I). The *Dynamic model* updates the probabilities after each concept, to minimize the expected description length of the concepts learnt so far. We include in this model the \oplus rule *a priori* in the language, but with vanishing probability of being used. Changes in this probability can be analogously interpreted as the probability that a rational agent without the compiled symbol *a priori* decides to add the compiled expression as a new primitive into her language.

A. Static model

Under the LoT assumption, given a concept C (e.g., Fig. 2), the probability that an agent uses formula φ to explain this concept is defined by Bayes theorem:

$$P(\varphi | C) \propto P(C | \varphi)P(\varphi).$$

The likelihood $P(C | \varphi)$ of a logical statement φ can be simply defined as 1 if $\llbracket \varphi \rrbracket = C$ and 0 otherwise. In other words, for any given concept, only explanations that describe this concept are considered as possible explanations. The likelihood term has been defined more flexibly in the literature [8,15], allowing for mislabeled elements. We keep this simpler definition to reduce the number of free parameters of the model, as we do not intend to account for mislabeling errors in our experiment.

The prior $P(\varphi)$ is defined by augmenting the context-free grammars shown in Fig. 1 into a probabilistic context-free grammars (PCFG). In the PCFG, each rule has associated a parameter indicating the probability of using that rule. A PCFG can be used to produce logical statements similar to a CFG. Each nonterminal remaining in the statement is expanded using a rule of the PCFG with probability proportional to that rule's associated parameter, until no nonterminals remain in the statement.

We assume that the probability that a subject uses formula φ to explain concept C is proportional to the posterior $P(\varphi | C)$, and the subjective difficulty d_C of a concept C to a participant is proportional to the length of the formula that the participant is using to explain that concept. However, there is no way to know directly which internal formula φ the participant is using (and therefore we do not know $|\varphi|$).

Hence, the most parsimonious approach is to consider the entire posterior distribution $\mathbf{P}(\varphi | \mathcal{C})$ over possible formulas.²

Given a concept \mathcal{C} , the expected length $E_{\mathcal{C}}$ of the formulas used by the participant is simply

$$E_{\mathcal{C}} = \sum_{\llbracket \varphi \rrbracket = \mathcal{C}} |\varphi| P(\varphi | \mathcal{C}), \quad (1)$$

where the sum is over all formulas φ compatible with \mathcal{C} . We define the difficulty $d_{\mathcal{C}}$ of a concept experienced by the participant as

$$d_{\mathcal{C}} \propto E_{\mathcal{C}} + \alpha N_{\mathcal{C}},$$

where we added a term that accounts for the cardinality of the concept: $N_{\mathcal{C}}$ is the cardinality of the concept or its complement, the one being smaller, i.e., $N_{\mathcal{C}} = \min\{\#\mathcal{C}, \#\bar{\mathcal{C}}\}$ (e.g., $N_{\mathcal{C}} = 4$ for the concept \mathcal{C} of Fig. 2), and α is a free parameter fitted globally for all concepts and participants to its maximum likelihood value of 0.9. In this way, we remove the asymmetry between positive and negative examples, while accounting for the toil taken by considering a larger number of items simultaneously.

In practice, to approximate $E_{\mathcal{C}}$ for each concept \mathcal{C} , we calculated the posterior probability $P(\varphi | \mathcal{C})$ of all compatible formulas φ s up to size 19 with $P(\varphi | \mathcal{C})$ and then use (1). Since the space of all possible φ s grows exponentially with $|\varphi|$, normative procedures for estimating $P(\varphi | \mathcal{C})$ in this space involve stochastic search algorithms. However, in our case, we were able to exhaustively enumerate and calculate the posterior probability of *all* formulas generated by the PCFG up to a sufficiently high size M such that all formulas with $|\varphi| > M$ have vanishing probabilities when compared to shorter compatible formulas for the current concept (because the prior $P(\varphi)$ decreases exponentially with the size of the formula).

B. Dynamic model

Up to this point, we assumed that, given a concept \mathcal{C} , the posterior distribution over formulas $P(\varphi | \mathcal{C})$ was independent of the other concepts presented in the sequence. However, if the LoT (i.e., the PCFG) updates with experience, then the prior $P(\varphi)$ in $P(\varphi | \mathcal{C})$ will change, and so will $E_{\mathcal{C}}$ in Eq. (1) and finally the subjective difficulty $d_{\mathcal{C}}$. Therefore, $d_{\mathcal{C}}$ will depend on the sequence of concepts that were previously presented to the participant.

In other words, since now $P(\varphi)$ depends on the sequence of concepts experienced by the participant, instead of $P(\varphi | \mathcal{C})$, we have

$$P(\varphi | \mathcal{C}^t, \dots, \mathcal{C}^1) \propto P(\mathcal{C}^t | \varphi) P(\varphi | \mathcal{C}^1, \dots, \mathcal{C}^{t-1}),$$

where \mathcal{C}^t is the concept presented at trial t , and $P(\varphi | \mathcal{C}^1, \dots, \mathcal{C}^{t-1})$ depends on the state of the PCFG at trial t , which in turn depends on how the PCFG gets updated from trial to trial.

²This is equivalent to the Sampling Hypothesis described in Ref. [18], by which participants represent distributions through samples. Similar results are obtained if each participant carries entire probability distributions.

Intuitively, the update process increases the probability of using a certain rule in the PCFG accordingly to how useful this rule was to compress compatible formulas for the concepts previously learned in the same domain. Specifically, we model the update process in a normative manner: the probability of using a rule of the PCFG at trial t is equal to the Bayesian posterior probability that this rule will enable the learner to find compressed explanations at trial t , according to how useful it was to compress explanations in trials $1, \dots, t-1$.

To formalize the update of the PCFG, we define $P(\varphi)$ similarly to Ref. [15]. Specifically, the prior probability of a logical statement at trial t in the concept sequence uses a single Dirichlet-multinomial for the set of rule expansions. The Dirichlet is parameterized by a set of positive real numbers D_i^t , one for each rule i in the PCFG, which in turn determine the probability of using rule i at trial t : a higher D_i indicates a higher probability of using rule i .

The prior is specified by the set Dirichlet parameters \mathbf{D}^0 with which we start the experiment (\mathbf{D}^0 represents a vector containing the prior parameters of all rules in the grammar at trial 0). In our experiment, we set the prior Dirichlet parameters of all rules equal to 1, and the parameter of the rule that expands the target operator to a value several orders of magnitude smaller ($\approx 10^{-4}$). This means that the target operator was practically absent at the beginning of the experiment, but it was technically possible to “learn it” by increasing its probability as the experiment developed.

Under the Dirichlet model, the prior $P(\varphi | \mathcal{C}^1, \dots, \mathcal{C}^{t-1})$ can be rewritten using the Dirichlet parameters as $P(\varphi | \mathbf{D}^t)$. Therefore, to know how $P(\varphi | \mathcal{C})$ updates from trial to trial, we only need to know how \mathbf{D} updates from trial to trial.

The Dirichlet parameter of rule i at trial $t+1$ is equal to its parameter at trial t plus the amount of times the production i was used in generating all formulas compatible with the concept at trial t (we note $M_i(\varphi)$ as the number of times that rule i is used in generating formula φ), weighted by each formula’s posterior probability at trial t :

$$D_i^{t+1} = D_i^t + \sum_{\llbracket \varphi \rrbracket = \mathcal{C}^t} P(\varphi | \mathbf{D}^t) M_i(\varphi). \quad (2)$$

This Bayesian learning mechanism increases the probability of using rules that allow concepts to be succinctly described. This happens because these formulas have higher probability $P(\varphi | \mathbf{D})$ than longer formulas, so the Dirichlet parameters of the rules that build these formulas increase more strongly than those of the rules that build longer formulas.

VI. RESULTS

The Bayesian agent that minimizes the expected complexity of future concepts by optimally adapting its LoT to the inferred structure of the task accurately captures the dynamics of human learning across concepts. If we did not allow the model to update the probability of the operators after each concept, and particularly the compiled operator \oplus , then the control group and the target group would be indistinguishable to the model as it would predict equal average formula length for both groups (see Fig. 3, static model). Instead, as shown in Fig. 4, by adjusting the prior probabilities based on concept exposure the dynamic model is able to capture learning time

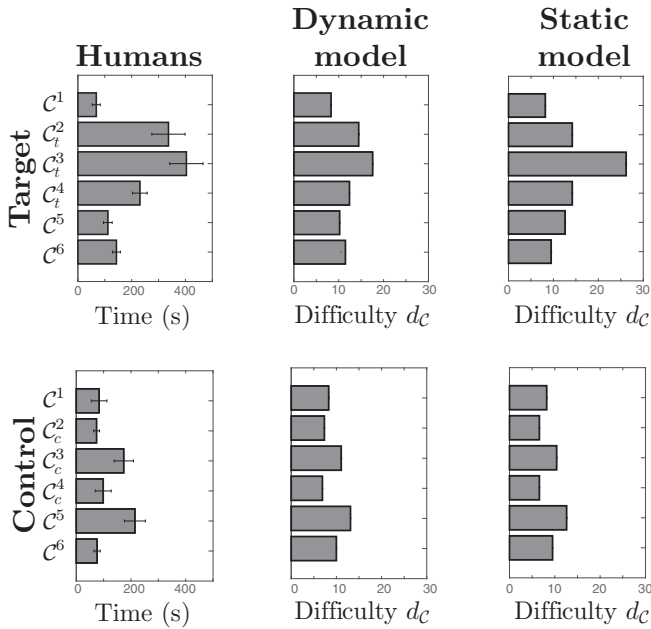


FIG. 4. Learning times and model predictions for target and control groups (see Table I for concept details). The predicted difficulties of each model were calculated using d_c . Error bars are s.e.m.

patterns in the target groups ($R^2 = 0.96$ compared to $R^2 = 0.73$ for the static model). Expectedly, both models perform similarly in the control groups as they were designed to not encourage the use of any particular operator ($R^2 = 0.72$; $R^2 = 0.71$ for the static model). The impact of the learning capability of the model is most evident in the target group concept sequence, which was designed to this effect. If the structure of the concepts does not bias the LoT primitives one way or the other, then it is expected that a static model will provide a reasonable fit. However, it is difficult to tell *a priori* how unbiased a set of concepts really is, so experiments relying on repeated concept exposure should always take between-concept learning into account.

Allowing the model to constantly update its beliefs from concept to concept is a requisite to capture human learning times. We now explain how the pattern of subjective difficulties in Fig. 4 emerged in the *dynamic model*. In this scenario, learning for the model is formalized by the update of rule parameters from concept t to concept $t + 1$ according to Eq. (2). In Fig. 5 we show how this learning takes place in the concept sequence for the target group. There are mainly




FIG. 5. Evolution of Dirichlet parameters of different rules after each concept experienced by the target group.

two competing formulas when C_t^2 is presented: $x_i \oplus x_j$ and $(x_i \wedge \neg x_j) \vee (\neg x_i \wedge x_j)$. Given the low *a priori* value of the parameter of the \oplus rule, the posterior of the formulas of type $(x_i \wedge \neg x_j) \vee (\neg x_i \wedge x_j)$, which do not use the \oplus operator, is much higher than the posterior of $x_i \oplus x_j$. Therefore, in Fig. 4 we see a large predicted difficulty by the dynamic model for this concept [since the posterior lies mainly over these longer formulas without \oplus , see Eq. (1)].

However, the little increment in the \oplus rule after C_t^2 (see Fig. 5) is sufficient for making the formula $x_k \oplus x_l$ to have higher relative posterior in the next concepts, making the increment in the parameter of the \oplus rule much greater than before. Additionally, the difficulty inferred by the model is much smaller the second time the concept is presented (compare C_t^4 and C_t^2 concepts in Fig. 4), since now the posterior is more evenly distributed between long (without \oplus) and short (with \oplus) formulas [see Eq. (1)]. Finally, when the concept C^5 is presented, the learner has completely compiled the \oplus rule into her language, ascribing the formulas that use the \oplus operator a much higher posterior probability relative to the long formulas that do not use the \oplus operator. Therefore, the inferred difficulty for C^5 is much smaller than those describing previous concepts, almost as simple as concept C^1 (see Fig. 4).

Finally, the strong \oplus acquired by the target group increases the difficulty of C^6 relative to the control group (see Fig. 3). This occurs because there are several formulas of length 9 that use the \oplus operator (around 6000), significantly increasing the expected difficulty of the concept [see Eq. (1)]. For the control group, the posterior probability of these formulas is very low, causing a smaller increase in the expected difficulty.

The previous results point to a competition between different rules in the grammar. In our model, competition between \oplus and the other operators is modulated by the initial relative value of the Dirichlet prior of the \oplus rule, and the overall magnitude of the priors of all rules. The initial \oplus prior measures how useful \oplus should be (relative to the other rules) to increase the likelihood of using it in the future. If the \oplus prior is too low relative to the priors of other rules, then formulas with \oplus must be much shorter than formulas without \oplus in order for them to have appreciable posterior and increase the \oplus parameter in Eq. (2). In our experiment, if the prior is smaller 10^{-12} (and 1 for all other rules), then the predictions of the dynamic and static model for the target group are approximately equal: the advantage of using \oplus in the target concepts is not enough to increase the likelihood of using \oplus . On the other hand, if the \oplus prior is too high, we cannot model the high difficulty of C_t^2 for the target group and the high difficulty of C^5 for the control group. For example, if the \oplus prior is higher than 0.05 (and 1 for all other rules), the difficulty of C_t^2 and C_t^4 are approximately equal (corresponding to the short formula with \oplus) and also the difficulties of C^5 for control and target groups.

The other free parameter that modulates competition is the overall magnitude of the Dirichlet priors, which determines how many times an efficient rule should be encountered before incorporating it. If the magnitude is too high, then observing a useful rule does not significantly change its Dirichlet parameter relative to the others, eliminating from the model the rapid rule acquisition clearly showed by participants. This happens because in Eq. (2) the magnitude of the updates from t to $t + 1$

are at most of order M , the number of times that operators appear in formulas with high posterior. In our experiment, if all rules have prior equal to 1 and \oplus has 1/1000 we get similar results to the ones in Fig. 4, but if all rules have prior equal to 10 000 and \oplus has 10 the additions to the \oplus parameter are insignificant, so the dynamic and static models make the same predictions for the target group.

In our model a large enough exposure to a concepts will increase the Dirichlet parameters without bounds, progressively decreasing learning flexibility. Although our experiment is not long enough to test it, such inflexibility is very unlikely to be true. For example, in the LoT fitting experiment from Ref. [8] they found that human Dirichlet priors for most propositional operators are between 0.3 and 3, instead of orders of magnitude higher [as expected by Eq. (2) after exposure to a large number of concepts]. Therefore, a more complete model of lifelong language acquisition should include an extra normalization or forgetting parameter that decreases the overall magnitude of the Dirichlet parameters, preserving the high learning flexibility that we observed in our experiment.

VII. DISCUSSION

We measured the subjective difficulty that participants experience when learning a sequence of concepts. To explain this subjective difficulty, we resource to propositional logic as a base description language. In the target group we experimented with concepts which can be succinctly described in the base language *that also contains an extra operator* \oplus for exclusive disjunction but that needed necessarily longer descriptions over the base language (where this operator is absent). On the contrary, the control group is exposed to concepts where \oplus does not help to achieve succinctness.

Learning times are consistent with the hypothesis that participants in the target group smoothly adopt the \oplus as a new primitive of their LoT to absorb the concepts they have been exposed to, with no more incentive than decreasing the expected complexity of future concepts. We do not claim that participants have learned the \oplus operator defined by any specific formula using the previous operators, however, their LoT seems to have constructed an operation that matches the semantics of the exclusive or to compress such patterns of data and identify them more efficiently.

Here, we focus on transfer learning effects when learning sequential concepts that share the same hierarchical structure. We acknowledge, however, that several other transfer learning effects are present in human sequential logical concept learning, such as when subsequent concepts differ in the relevant variables (e.g., color lights in our experiment) [19], when changing the relevant variables in subsequent exclusive disjunctions [20], or when two categories are learned in an interleaved or a focused manner [21]. However, unlike superficial knowledge about the task (like the frequency of

appearance of different symbols and logical operators in the concept sequence), identifying the latent hierarchical structure of concepts have extremely important computational consequences: it allows for exponentially less complex representations [22,23], maximizing the expected value of future computations within resource-bounded constraints [24]. In our task, to focus primarily on the learning process of the \oplus structure, we randomize variables in each trial, such that other kinds of transitions are averaged out across participants.

Most LoT studies provide a language that is fixed once trained or inferred over a specific data. We claim that when a specific language beats a second one at fitting some experimental data, what we may be seeing is an effect of prior experience (including from the experiment itself), more than an intrinsic feature of the LoT. This leads to a fundamental difficulty in trying to experimentally uncover what the actual human symbolic substrate of thought is. Experimental results have shown for instance that a grammar with *and*, *or*, and *not* better explains Boolean concept learning than one with *nand*, despite both being expressively equivalent [8]. In our view, this cannot be taken to mean anything more than that in the current state of affairs of the world, the *nand* operator is not very useful for compressing information. We have shown that participants can rapidly compile new expressions in their LoT if they begin to be useful, which emphasizes that one cannot simply ignore the order in which concepts are presented to the participant when studying aspects of the LoT.

When Fodor proposed the language of thought hypothesis [4], what he had in mind was a symbolic system we all came equipped with from birth. Stating that this language is in fact always flexible might seem in outright contradiction with Fodor's original idea. In fact, what studies in the LoT literature (including this one) are probably probing is one among many languages in a hierarchy of increasing abstraction. As we progress in life, we find some conceptual summaries useful, and compiled them in a more abstract token. It is even likely that there is no proper hierarchy with sharply defined boundaries between levels, but instead a less organized progression of concepts of increasing abstraction, with thought progressing seamlessly using constructs at different levels.

VIII. CONCLUSION

We defined a model to measure the subjective difficulty of learning a sequence of concepts. The model updates the grammar production probabilities between concepts and predicts difficulty as the size of compatible formulas weighted by their posterior probability. This learning mechanism allows to simulate the emergence of a new primitive in the language, as it becomes useful to encode the concepts presented so far. The predicted difficulties strongly resembles the pattern of human learning times in a sequence of concepts that required the \oplus operator to be efficiently represented.

[1] N. Chomsky, *Knowledge of Language: Its Nature, Origin, and Use* (Greenwood Publishing Group, Santa Barbara, CA, (1986).

[2] N. Chomsky *et al.*, On cognitive structures and their development: A reply to Piaget, in *Philosophy of Mind: Classical Problems/Contemporary Issues*,

- edited by B. Beakley and P. Ludlow (MIT Press, Cambridge, MA, 2006), p. 751.
- [3] A. Newell, Physical symbol systems, *Cogn. Sci.* **4**, 135 (1980).
- [4] J. A. Fodor, *The Language of Thought* (Harvard University Press, Boston, MA, 1975), Vol. 5.
- [5] H. Abelson, N. Goodman, and L. Rudolph, *Logo Manual* (MIT Press, Cambridge, MA, 1974).
- [6] A. N. Kolmogorov, Three approaches to the quantitative definition of information, *Int. J. Comput. Math.* **2**, 157 (1968).
- [7] J. Feldman, The simplicity principle in human concept learning, *Curr. Direct. Psychol. Sci.* **12**, 227 (2003).
- [8] S. T. Piantadosi, J. B. Tenenbaum, and N. D. Goodman, The logical primitives of thought: Empirical foundations for compositional cognitive models, *Psychol. Rev.* **123**, 392 (2016).
- [9] E. L. Leeuwenberg, A perceptual coding language for visual and auditory patterns, *Am. J. Psychol.* **84**, 307 (1971).
- [10] M. Amalric, L. Wang, P. Pica, S. Figueira, M. Sigman, and S. Dehaene, The language of geometry: Fast comprehension of geometrical primitives and rules in human adults and preschoolers, *PLoS Comput. Biol.* **13**, e1005273 (2017).
- [11] S. Romano, A. Salles, M. Amalric, S. Dehaene, M. Sigman, and S. Figueira, Bayesian validation of grammar productions for the language of thought, *PLoS One* **13**, e0200420 (2018).
- [12] G. Lupyan, D. H. Rakison, and J. L. McClelland, Language is not just for talking: Redundant labels facilitate learning of novel categories, *Psychol. Sci.* **18**, 1077 (2007).
- [13] M. L. Minsky, *Computation: Finite and Infinite Machines* (Prentice-Hall, Upper Saddle River, NJ, 1967).
- [14] G. L. Murphy, Comprehending complex concepts, *Cogn. Sci.* **12**, 529 (1988).
- [15] N. D. Goodman, J. B. Tenenbaum, J. Feldman, and T. L. Griffiths, A rational analysis of rule-based concept learning, *Cogn. Sci.* **32**, 108 (2008).
- [16] C. Kemp, Exploring the conceptual universe, *Psychol. Rev.* **119**, 685 (2012).
- [17] <https://figshare.com/s/04d338adbbc4b1e83bf0>.
- [18] S. Denison, E. Bonawitz, A. Gopnik, and T. L. Griffiths, Rational variability in children's causal inferences: The sampling hypothesis, *Cognition* **126**, 285 (2013).
- [19] M. R. Blair, M. R. Watson, R. C. Walshe, and F. Maj, Extremely selective attention: Eye-tracking studies of the dynamic allocation of attention to stimulus features in categorization, *J. Exper. Psychol.: Learn. Mem. Cogn.* **35**, 1196 (2009).
- [20] J. K. Kruschke, Dimensional relevance shifts in category learning, *Connect. Sci.* **8**, 225 (1996).
- [21] P. F. Carvalho and R. L. Goldstone, Putting category learning in order: Category structure and temporal arrangement affect the benefit of interleaved over blocked study, *Mem. Cogn.* **42**, 481 (2014).
- [22] Y. Bengio, A. Courville, and P. Vincent, Representation learning: A review and new perspectives, *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 1798 (2013).
- [23] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum, Human-level concept learning through probabilistic program induction, *Science* **350**, 1332 (2015).
- [24] S. J. Gershman, E. J. Horvitz, and J. B. Tenenbaum, Computational rationality: A converging paradigm for intelligence in brains, minds, and machines, *Science* **349**, 273 (2015).