# Reinforcement and inference in cross-situational word learning

## Paulo F. C. Tilles and José F. Fontanari *

*Departamento de Física e Informática, Instituto de Física de São Carlos, Universidade de São Paulo, São Carlos, Brazil*

Cross-situational word learning is based on the notion that a learner can determine the referent of a word by finding something in common across many observed uses of that word. Here we propose an adaptive learning algorithm that contains a parameter that controls the strength of the reinforcement applied to associations between concurrent words and referents, and a parameter that regulates inference, which includes built-in biases, such as mutual exclusivity, and information of past learning events. By adjusting these parameters so that the model predictions agree with data from representative experiments on cross-situational word learning, we were able to explain the learning strategies adopted by the participants of those experiments in terms of a trade-off between reinforcement and inference. These strategies can vary wildly depending on the conditions of the experiments. For instance, for fast mapping experiments (i.e., the correct referent could, in principle, be inferred in a single observation) inference is prevalent, whereas for segregated contextual diversity experiments (i.e., the referents are separated in groups and are exhibited with members of their groups only) reinforcement is predominant. Other experiments are explained with more balanced doses of reinforcement and inference.

Keywords: statistical learning, word learning, cross-situational learning, associative learning, mutual exclusivity

## 1. INTRODUCTION

A desirable goal of a psychological theory is to offer explanations grounded on elementary principles to the data available from psychology experiments (Newell, 1994). Although most of these quantitative psychological data are related to mental chronometry and memory accuracy, recent explorations on the human performance to acquire an artificial lexicon in controlled laboratory conditions have paved the way to the understanding of the learning strategies humans use to infer a word-object mapping (Yu and Smith, 2007; Kachergis et al., 2009; Smith et al., 2011; Kachergis et al., 2012; Yu and Smith, 2012a). These experiments are based on the cross-situational word-learning paradigm which avers that a learner can determine the meaning of a word by finding something in common across all observed uses of that word (Gleitman, 1990; Pinker, 1990). In that sense, learning takes place through the statistical sampling of the contexts in which a word appears in accord with the classical associationist stance of Hume and Locke that the mechanism of word learning is sensitivity to covariation: if two events occur at the same time, they become associated (Bloom, 2000).

In a typical cross-situational word-learning experiment, participants are exposed repeatedly to multiple unfamiliar objects concomitantly with multiple spoken pseudo-words, such that a word and its correct referent (object) always appear together on a learning trial. Different trials exhibiting distinct word-object pairs will eventually allow the disambiguation of the word-object associations and the learning of the correct mapping (Yu and Smith, 2007). However, it is questionable whether this scenario is suitable to describe the actual word learning process by children even in the unambiguous situation where the single novel object is followed by the utterance of its corresponding pseudo-word. In fact, it was shown that young children will only make the connection between the object and the word provided they have a reason to believe that they are in presence of an act of naming and for this the speaker has to be present (Baldwin et al., 1996; Bloom, 2000; Waxman and Gelman, 2009). Adults could learn those associations either because they were previously instructed by the experimenter that they would be learning which words go with which objects or because they could infer that the disembodied voice is an act of naming by a concealed person. Although there have been claims that cross-situational statistical learning is part of the repertoire of young word learners (Yu and Smith, 2008), the effect of individual differences in attention and vocabulary development of the infants complicates considerably this issue which is still a matter for debate (Yu and Smith, 2012b; Smith and Yu, 2013).

There are several other alternative or complementary approaches to the statistical learning formulation of language acquisition considered in this paper. For instance, the social-pragmatic hypothesis claims that the child makes the connections between words and their referents by understanding the referential intentions of others. This approach, which seems to be originally due to Augustine, implies that children use intuitive psychology to "read" the adults' minds (Bloom, 2000). A more recent approach that explores the grounding of language in perception and action has been proved effective in the design of

linguistic capabilities in humanoid cognitive robots (Cangelosi et al., 2007; Cangelosi, 2010; Pezzulo et al., 2013) as well as in the support of word learning by toddlers through the stabilization of their attention on the selected object (Yu and Smith, 2012b). In contrast with the unsupervised cross-situational learning scheme, the scenario known as operant conditioning involves the active participation of the agents in the learning process, with exchange of non-linguistic cues to provide feedback on the learner inferences. This supervised learning scheme has been applied to the design of a system for communication by autonomous robots in the Talking Heads experiments (Steels, 2003). We note that a comparison between the cross-situational and operant conditioning learning schemes indicates that they perform similarly in the limit of very large lexicon sizes (Fontanari and Cangelosi, 2011).

As our goal is to interpret the learning performance of adults using a few plausible reasoning tenets, here we assume that in order to learn a word-object mapping within the cross-situational word-learning scenario the learner should be able to (i) recall at least a fraction of the word-object pairings that appeared in the learning trials, (ii) register both co-occurrences and non-co-occurrences of words and objects and (iii) apply the mutual exclusivity principle which favors the association of novel words to novel objects (Markman and Wachtel, 1988). Of course, we note that a hypothetical learner could achieve cross-situational learning solely by registering and recalling co-occurrences of words and objects without carrying out any inferential reasoning (Blythe et al., 2010; Tilles and Fontanari, 2012a), but we find it implausible that human learners would not reap the benefits (e.g., fast mapping) of employing mutual exclusivity (Vogt, 2012; Reisenauer et al., 2013).

In this paper we offer an adaptive learning algorithm that comprises two parameters which regulate the associative reinforcement of pairings between concurrent words and objects, and the non-associative inference process that handles built-in biases (e.g., mutual exclusivity) as well as information of past learning events. By setting the values of these parameters so as to fit a representative selection of experimental data presented in Kachergis et al. (2009, 2012) we are able to identify and explain the learning strategies adopted by the participants of those experiments in terms of a trade-off between reinforcement and inference.

## 2. CROSS-SITUATIONAL LEARNING SCENARIO

We assume there are $N$ objects $o_1, \ldots, o_N$, $N$ words $w_1, \ldots, w_N$ and a one-to-one mapping between words and objects represented by the set $\Gamma = \{(w_1, o_1), \ldots, (w_N, o_N)\}$. At each learning trial, $C$ word-object pairs are selected from $\Gamma$ and presented to the learner without providing any clue on which word goes with which object. For instance, pictures of the $C$ objects are displayed in a slide while $C$ pseudo-words are spoken sequentially such that their spatial and temporal arrangements do not give away the correct word-object associations (Yu and Smith, 2007; Kachergis et al., 2009). We refer to the subset of words and their referents (objects) presented to the learner in a learning trial as the context $\Omega = \{w_1, o_1, w_2, o_2, \ldots, w_C, o_C\}$. The context size $C$ is then a measure of the within-trial ambiguity, i.e., the number of co-occurring

word-object pairs per learning trial. The selection procedure from the set $\Gamma$, which may favor some particular subsets of word-object pairs, determines the different experimental setups discussed in this paper. Although each individual trial is highly ambiguous, repetition of trials with partially overlapping contexts should in principle allow the learning of the $N$ word-object associations.

After the training stage is completed, which typically comprises about two dozen trials, the learning accuracy is measured by instructing the learner to pick the object among the $N$ objects on display which the learner thinks is associated to a particular target word. The test is repeated for all $N$ words and the average learning accuracy calculated as the fraction of correct guesses (Kachergis et al., 2009).

This cross-situational learning scenario does not account for the presence of noise, such as the effect of out-of-context words. This situation can be modeled by assuming that there is a certain probability (noise) that the referent of one of the spoken words is not part of the context (so that word can be said to be out of context). Although theoretical analysis shows that there is a maximum noise intensity beyond which statistical learning is unattainable (Tilles and Fontanari, 2012b), as yet no experiment was carried out to verify the existence of this threshold phenomenon on the learning performance of human subjects.

## 3. MODEL

We model learning as a change in the confidence with which the algorithm (or, for simplicity, the learner) associates the word $w_i$ to an object $o_j$ that results from the observation and analysis of the contexts presented in the learning trials. More to the point, this confidence is represented by the probability $P_t(w_i, o_j)$ that $w_i$ is associated to $o_j$ at learning trial $t$. This probability is normalized such that $\sum_{o_j} P_t(w_i, o_j) = 1$ for all $w_i$ and $t > 0$, which then implies that when the word $w_i$ is presented to the learner in the testing stage the learning accuracy is given simply by $P_t(w_i, o_i)$. In addition, we assume that $P_t(w_i, o_j)$ contains information presented in the learning trials up to and including trial $t$ only.

If at learning trial $t$ the learner observes the context $\Omega_t = \{w_1, o_1, w_2, o_2, \ldots, w_C, o_C\}$ then it can infer the existence of two other informative sets. First, the set of the words (and their referents) that appear for the first time at trial $t$, which we denote by $\tilde{\Omega}_t = \{\tilde{w}_1, \tilde{o}_1, \tilde{w}_2, \tilde{o}_2, \ldots, \tilde{w}_{\tilde{C}}, \tilde{o}_{\tilde{C}_t}\}$. Clearly, $\tilde{\Omega}_t \subseteq \Omega_t$ and $\tilde{C}_t \leq C$. Second, the set of words (and their referents) that do not appear in $\Omega_t$ but that have already appeared in the previous trials, $\bar{\Omega}_t = \{\bar{w}_1, \bar{o}_1, \ldots, \bar{w}_{N_t - C}, \bar{o}_{N_t - C}\}$ where $N_t$ is the total number of different words that appeared in contexts up to and including trial $t$. Clearly, $\bar{\Omega}_t \cap \Omega_t = \varnothing$. The update rule of the confidences $P_t(w_i, o_j)$ depends on which of these three sets the word $w_i$ and the object $o_j$ belong to (if $i \neq j$ they may belong to different sets). In fact, our learning algorithm comprises a parameter $\chi \in [0, 1]$ that measures the associative reinforcement capacity and applies only to known words that appear in the current context, and a parameter $\beta \in [0, 1]$ that measures the inference capacity and applies either to known words that do not appear in the current context or to new words in the current context. Before the

experiment begins ($t = 0$) we set $P_0\left(w_i, o_j\right) = 0$ for all words $w_i$ and objects $o_j$. Next we describe how the confidences are updated following the sequential presentation of contexts.

In the first trial ($t = 1$) all words are new ($\tilde{C}_1 = N_1 = C$), so we set

$$P_1\left(\tilde{w}_i, \tilde{o}_j\right) = \frac{1}{C} \qquad (1)$$

for $\tilde{w}_i, \tilde{o}_j \in \tilde{\Omega} = \Omega$. In the second or in an arbitrary trial $t$ we expect to observe contexts exhibiting both novel and repeated words. Novel words must go through an inference preprocessing stage before the reinforcement procedure can be applied to them. This is so because if $\tilde{w}_i$ appears for the first time at trial $t$ then $P_{t-1}\left(\tilde{w}_i, o_j\right) = 0$ for all objects $o_j$ and since the reinforcement is proportional to $P_{t-1}\left(\tilde{w}_i, o_j\right)$ the confidences associated to $\tilde{w}_i$ would never be updated (see Equation 5 and the explanation thereafter). Thus, when a novel word $\tilde{w}_i$ appear at trial $t \geq 2$, we redefine its confidence values at the previous trial (originally set to zero) as

$$P_{t-1}\left(\tilde{w}_i, \tilde{o}_j\right) = \frac{\beta}{\tilde{C}_t} + \frac{1-\beta}{N_{t-1} + \tilde{C}_t}, \qquad (2)$$

$$P_{t-1}\left(\tilde{w}_i, o_j\right) = \frac{1-\beta}{N_{t-1} + \tilde{C}_t}, \qquad (3)$$

$$P_{t-1}\left(\tilde{w}_i, \bar{o}_j\right) = \frac{1-\beta}{N_{t-1} + \tilde{C}_t}. \qquad (4)$$

On the one hand, setting the inference parameter $\beta$ to its maximum value $\beta = 1$ enforces the mutual exclusivity principle which requires that the new word $\tilde{w}_i$ be associated with equal probability to the $\tilde{C}_t$ new objects $\tilde{o}_j$ in the current context. Hence in the case $\tilde{C}_t = 1$ the meaning of the new word would be inferred in a single presentation. On the other hand, for $\beta = 0$ the new word is associated with equal probability to all objects already seen up to and including trial $t$, i.e., $N_t = N_{t-1} + \tilde{C}_t$. Intermediate values of $\beta$ describe a situation of imperfect inference. Note that using Equations 2–4 we can easily verify that $\sum_{\tilde{o}_j} P_{t-1}\left(\tilde{w}_i, \tilde{o}_j\right) + \sum_{o_j} P_{t-1}\left(\tilde{w}_i, o_j\right) + \sum_{\bar{o}_j} P_{t-1}\left(\tilde{w}_i, \bar{o}_j\right) = 1$, in accord with the normalization constraint.

Now we can focus on the update rule of the confidence $P_t\left(w_i, o_j\right)$ in the case both word $w_i$ and object $o_j$ appear in the context at trial $t$. The rule applies both to repeated and novel words, provided the confidences of the novel words are preprocessed according to Equations 2–4. In order to fulfill automatically the normalization condition for word $w_i$, the increase of the confidence $P_t\left(w_i, o_j\right)$ with $o_j \in \Omega_t$ must be compensated by the decrease of the confidences $P_t\left(w_i, \bar{o}_j\right)$ with $\bar{o}_j \in \bar{\Omega}_t$. This can be implemented by distributing evenly the total flux of probability out of the latter confidences, i.e., $\sum_{\bar{o}_j \in \bar{\Omega}_t} P_{t-1}\left(w_i, \bar{o}_j\right)$, over the confidences $P_t\left(w_i, o_j\right)$ with $o_j \in \Omega_t$. Hence the net gain of confidence on the association between $w_i$ and $o_j$ is given by

$$r_{t-1}\left(w_i, o_j\right) = \chi P_{t-1}\left(w_i, o_j\right) \frac{\sum_{\bar{o}_j \in \bar{\Omega}_t} P_{t-1}\left(w_i, \bar{o}_j\right)}{\sum_{o_j \in \Omega_t} P_{t-1}\left(w_i, o_j\right)} \qquad (5)$$

where, as mentioned before, the parameter $\chi \in [0, 1]$ measures the strength of the reinforcement process. Note that if both $o_j$ and $o_k$ appear in the context together with $w_i$ then the reinforcement procedure should not create any distinction between the associations $\left(w_i, o_j\right)$ and $(w_i, o_k)$. This result is achieved provided that the ratio of the confidence gains equals the ratio of the confidences before reinforcement, i.e., $r_{t-1}\left(w_i, o_j\right)/r_{t-1}(w_i, o_k) = P_{t-1}\left(w_i, o_j\right)/P_{t-1}(w_i, o_k)$. This is the reason that the reinforcement gain of a word-object association given by Equation 5 is proportional to the previous confidence on that association. The total increase in the confidences between $w_i$ and the objects that appear in the context, i.e., $\sum_{o_j \in \Omega_t} r_{t-1}\left(w_i, o_j\right)$, equals the product of $\chi$ and the total decrease in the confidences between $w_i$ and the objects that do not appear in the context, i.e., $\sum_{\bar{o}_j \in \bar{\Omega}_t} P_{t-1}\left(w_i, \bar{o}_j\right)$. So for $\chi = 1$ the confidences associated to objects absent from the context are fully transferred to the confidences associated to objects present in the context. Lower values of $\chi$ allows us to control the flow of confidence from objects in $\bar{\Omega}_t$ to objects in $\Omega_t$.

Most importantly, in order to implement the reinforcement process the learner should be able to gauge the relevance of the information about the previous trials, which is condensed on the confidence values $P_t\left(w_i, o_j\right)$. The gauging of this information is quantified by the word and trial dependent quantity $\alpha_t(w_i) \in [0, 1]$ that allows for the interpolation between the cases of maximum relevance ($\alpha_t(w_i) = 1$) and complete irrelevancy ($\alpha_t(w_i) = 0$) of the information stored in the confidences $P_t\left(w_i, o_j\right)$. In particular, we assume that the greater the certainty on the association between word $w_i$ and its referent, the more relevant that information is to the learner. A quantitative measure of the uncertainty associated to the confidences regarding word $w_i$ is given by the entropy

$$H_t(w_i) = -\sum_{o_j \in \Omega_t \cup \bar{\Omega}_t} P_t\left(w_i, o_j\right) \log\left[P_t\left(w_i, o_j\right)\right] \qquad (6)$$

whose maximum ($\log N_t$) is obtained by the uniform distribution $P_t\left(w_i, o_j\right) = 1/N_t$ for all $o_j \in \Omega_t \cup \bar{\Omega}_t$, and whose minimum (0) by $P_t\left(w_i, o_j\right) = 1$ and $P_t(w_i, o_k) = 0$ for $o_k \neq o_j$. So we define

$$\alpha_t(w_i) = \alpha_0 + (1 - \alpha_0)\left[1 - \frac{H_t(w_i)}{\log N_t}\right], \qquad (7)$$

where $\alpha_0 \in [0, 1]$ is a baseline information gauge factor corresponding to the maximum uncertainty about the referent of a target word.

Finally, recalling that at trial $t$ the learner has access to the sets $\Omega_t$, $\bar{\Omega}_t$ as well as to the confidences at trial $t - 1$ we write the update rule

$$P_t\left(w_i, o_j\right) = P_{t-1}\left(w_i, o_j\right) + \alpha_{t-1}(w_i)r_{t-1}\left(w_i, o_j\right)$$
$$+ \left[1 - \alpha_{t-1}(w_i)\right]\left[\frac{1}{N_t} - P_{t-1}\left(w_i, o_j\right)\right] \qquad (8)$$

for $w_i, o_j \in \Omega_t$. Note that if $\alpha_{t-1}(w_i) = 0$ the learner would associate word $w_i$ to all objects that have appeared up to and

including trial $t$ with equal probability. This situation happens only if $\alpha_0 = 0$ and if there is complete uncertainty about the referent of word $w_i$. Hence the quantity $\alpha_t(w_i)$ determines the extent to which the previous confidences on associations involving word $w_i$ influence the update of those confidences.

Now we consider the update rule for the confidence $P_t(w_i, \bar{o}_j)$ in the case that word $w_i$ appears in the context at trial $t$ but object $\bar{o}_j$ does not. (We recall that object $\bar{o}_j$ must have appeared in some previous trial.) According to the reasoning that led to Equation 5 this confidence must decrease by the amount $\chi P_{t-1}(w_i, \bar{o}_j)$ and so, taking into account the information gauge factor, we obtain

$$
P_t(w_i, \bar{o}_j) = P_{t-1}(w_i, \bar{o}_j) - \alpha_{t-1}(w_i)\,\chi P_{t-1}(w_i, \bar{o}_j)
$$
$$
+ [1 - \alpha_{t-1}(w_i)]\left[\frac{1}{N_t} - P_{t-1}(w_i, \bar{o}_j)\right] \quad (9)
$$

which can be easily seen to satisfy the normalization

$$
\sum_{o_j \in \Omega_t} P_t(w_i, o_j) + \sum_{\bar{o}_j \in \bar{\Omega}_t} P_t(w_i, \bar{o}_j) = 1. \quad (10)
$$

We focus now on the update rule for the confidence $P_t(\bar{w}_i, \bar{o}_j)$ with $\bar{w}_i, \bar{o}_j \in \bar{\Omega}_t$, i.e., both the word $\bar{w}_i$ and the object $\bar{o}_j$ are absent from the context shown at trial $t$, but they have already appeared, not necessarily together, in previous trials. A similar inference reasoning that led to the expressions for the preprocessing of new words would allow the learner to conclude that a word absent from the context should be associated to an object that is also absent from it. In that sense, confidence should flow from the associations between $\bar{w}_i$ and objects $o_j \in \Omega_t$ to the associations between $\bar{w}_i$ and objects $\bar{o}_j \in \bar{\Omega}_t$. Hence, ignoring the information gauge factor for the moment, the net gain to confidence $P_t(\bar{w}_i, \bar{o}_j)$ is given by

$$
\bar{r}_{t-1}(\bar{w}_i, \bar{o}_j) = \beta P_{t-1}(\bar{w}_i, \bar{o}_j) \frac{\sum_{o_j \in \Omega_t} P_{t-1}(\bar{w}_i, o_j)}{\sum_{\bar{o}_j \in \bar{\Omega}_t} P_{t-1}(\bar{w}_i, \bar{o}_j)}. \quad (11)
$$

The direct proportionality of this gain to $P_{t-1}(\bar{w}_i, \bar{o}_j)$ can be justified by an argument similar to that used to justify Equation 5 in the case of reinforcement. The information relevance issue is also handled in a similar manner so the desired update rule reads

$$
P_t(\bar{w}_i, \bar{o}_j) = P_{t-1}(\bar{w}_i, \bar{o}_j) + \alpha_{t-1}(\bar{w}_i)\,\bar{r}_{t-1}(\bar{w}_i, \bar{o}_j)
$$
$$
+ [1 - \alpha_{t-1}(\bar{w}_i)]\left[\frac{1}{N_t} - P_{t-1}(\bar{w}_i, \bar{o}_j)\right] \quad (12)
$$

for $\bar{w}_i, \bar{o}_j \in \bar{\Omega}_t$. To ensure normalization the confidence $P_t(\bar{w}_i, o_j)$ must decrease by an amount proportional to $\beta P_{t-1}(\bar{w}_i, o_j)$ so that

$$
P_t(\bar{w}_i, o_j) = P_{t-1}(\bar{w}_i, o_j) - \alpha_{t-1}(\bar{w}_i)\,\beta P_{t-1}(\bar{w}_i, o_j)
$$
$$
+ [1 - \alpha_{t-1}(\bar{w}_i)]\left[\frac{1}{N_t} - P_{t-1}(\bar{w}_i, o_j)\right] \quad (13)
$$

for $\bar{w}_i \in \bar{\Omega}_t$ and $o_j \in \Omega_t$. We can verify that prescriptions (12) and (13) satisfy the normalization

$$
\sum_{\bar{o}_j \in \bar{\Omega}_t} P_t(\bar{w}_i, \bar{o}_j) + \sum_{o_j \in \Omega_t} P_t(\bar{w}_i, o_j) = 1, \quad (14)
$$

as expected.

In summary, before any trial ($t = 0$) we set all confidence values to zero, i.e., $P_0(w_i, o_j) = 0$, and fix the values of the parameters $\alpha_0$, $\chi$ and $\beta$. In the first trial ($t = 1$) we set the confidences of the words and objects in $\Omega_1$ according to Equation (1), so we have the values of $P_1(w_i, o_j)$ for $w_i, o_j \in \Omega_1$. In the second trial, we separate the novel words $\tilde{w}_i \in \tilde{\Omega}_2$ and reset $P_1(\tilde{w}_i, o_j)$ with $o_i \in \Omega_2 \cup \bar{\Omega}_2$ according to Equations 2–4. Only then we calculate $\alpha_1(w_i)$ with $w_i \in \Omega_1 \cup \tilde{\Omega}_2$ using Equation (7). The confidences at trial $t = 2$ then follows from Equations (8), (9), (12), and (13). As before, in the third trial we separate the novel words $\tilde{w}_i \in \tilde{\Omega}_3$, reset $P_2(\tilde{w}_i, o_j)$ with $o_i \in \Omega_3 \cup \bar{\Omega}_3$ according to Equations 2–4, calculate $\alpha_2(w_i)$ with $w_i \in \Omega_1 \cup \Omega_2 \cup \tilde{\Omega}_3$ using Equation (7), and only then resume the evaluation of the confidences at trial $t = 3$. This procedure is repeated until the training stage is completed, say, at $t = t^*$. At this point, knowledge of the confidence values $P_{t^*}(w_i, o_j)$ allows us to answer any question posed in the testing stage.

Our model borrows many features from other proposed models of word learning (Siskind, 1996; Fontanari et al., 2009; Frank et al., 2009; Fazly et al., 2010; Kachergis et al., 2012). In particular, the entropy expression (6) was used by Kachergis et al. (2012) to allocate attention trial-by-trail to the associations presented in the contexts. Here we use that expression to quantify the uncertainty associated to the various confidences in order to determine the extent to which those confidences are updated on a learning trial. A distinctive feature of our model is the update of associations that are not in the current trial according to Equation (12). In particular, we note that whereas *ad hoc* normalization can only decrease the confidences on associations between words and objects that did not appear in the current context, our update rule can increase those associations as well. The extent of this update is weighted by the inference parameter $\beta$ and it allows the application of mutual exclusivity to associations that are not shown in the current context. In fact, the splitting of mental processes in two classes, namely, reinforcement processes that update associations in the current context and inference processes that update the other associations is the main thrust of our paper. In the next section we evaluate the adequacy of our model to describe a selection of cross-situational word-learning experiments carried out on adult subjects by Kachergis et al. (2009, 2012).

## 4. RESULTS

The cross-situational word-learning experiments of Kachergis et al. (2009, 2012) aimed to understand how word sampling frequency (i.e., number of trials in which a word appears), contextual diversity (i.e., the co-occurrence of distinct words or groups of words in the learning trials), within-trial ambiguity (i.e., the context size $C$), and fast-mapping of novel words affect the learning performance of adult subjects. In this section we compare

the performance of the algorithm described in the previous section with the performance of adult subjects reported in Kachergis et al. (2009, 2012). In particular, once the conditions of the training stage are specified, we carry out $10^4$ runs of our algorithm for fixed values of the three parameters $\alpha_0$, $\beta$, $\chi$, and then calculate the average accuracy at trial $t = t^*$ over all those runs for that parameter setting. Since the algorithm is deterministic, what changes in each run is the composition of the contexts at each learning trial. As our goal is to model the results of the experiments, we search the space of parameters to find the setting such that the performance of the algorithm matches that of humans within the error bars (i.e., one standard deviation) of the experiments.

## 4.1. WORD SAMPLING FREQUENCY

In these experiments the number of words (and objects) is $N = 18$ and the training stage totals $t^* = 27$ learning trials, with each trial comprising the presentation of 4 words together with their referents ($C = 4$). Following Kachergis et al. (2009), we investigate two conditions which differ with respect to the number of times a word is exhibited in the training stage. In the two-frequency condition, the 18 words are divided into two subsets of 9 words each. The words in the first subset appear 9 times and those in the second only 3 times. In the three-frequency condition, the 18 words are divided into three subsets of 6 words each. Words in the first subset appear 3 times, in the second, 6 times and in the third, 9 times. In these two conditions, the same word was not allowed to appear in two consecutive learning trials.
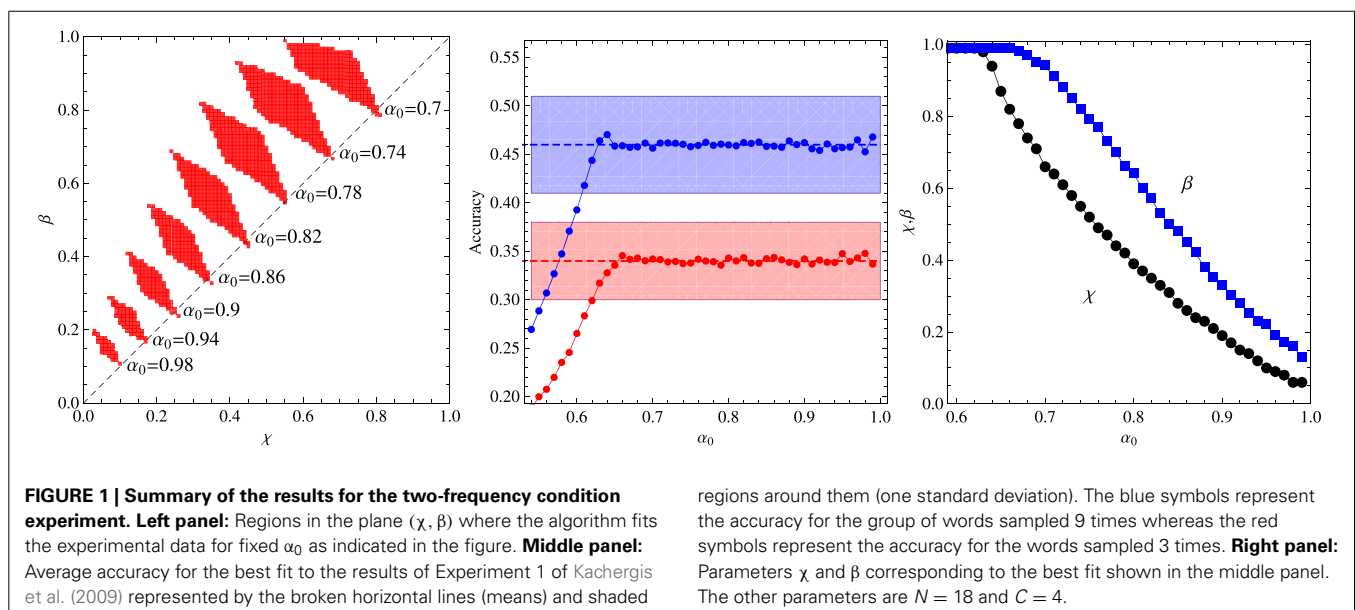
**Figures 1**, **2** summarize our main results for the two-frequency and three-frequency conditions, respectively. The left panels show the regions (shaded areas) in the ($\chi$, $\beta$) plane for fixed $\alpha_0$ where the algorithm describes the experimental data. We note that if those regions are located left to the diagonal $\chi = \beta$ then the inference process is dominant whereas if they are right to the diagonal then reinforcement is the dominant process. The middle panels show the accuracy of the best fit as function of the parameter $\alpha_0$
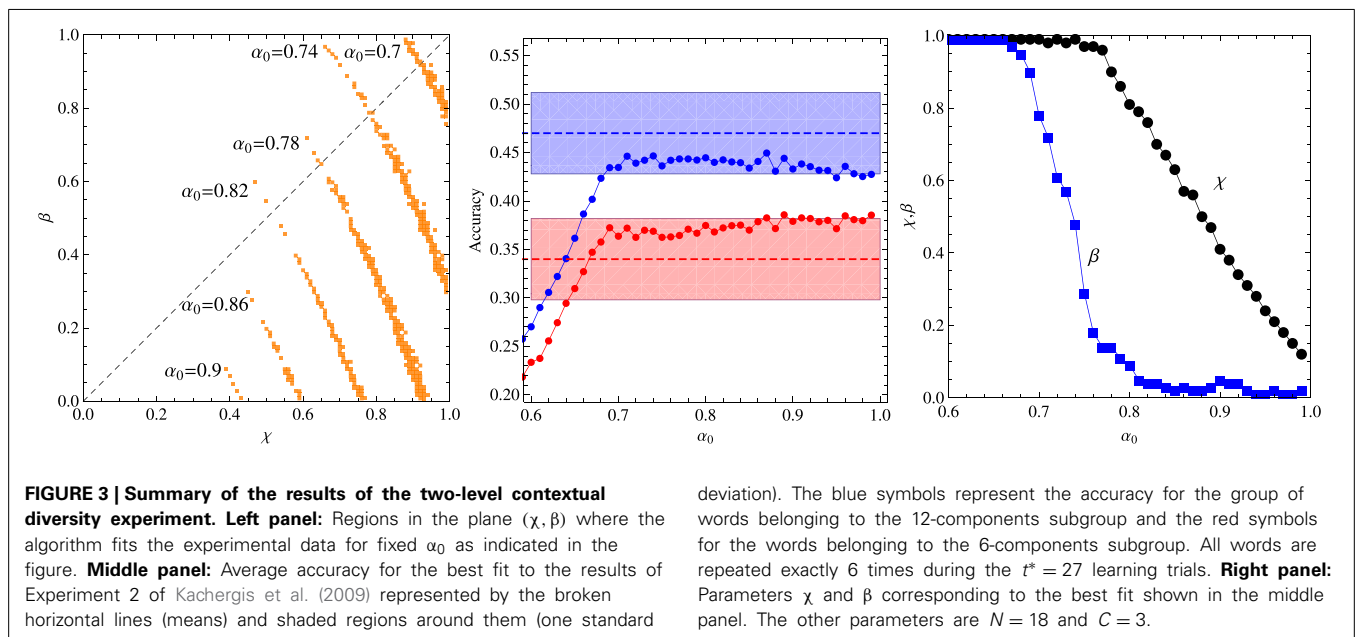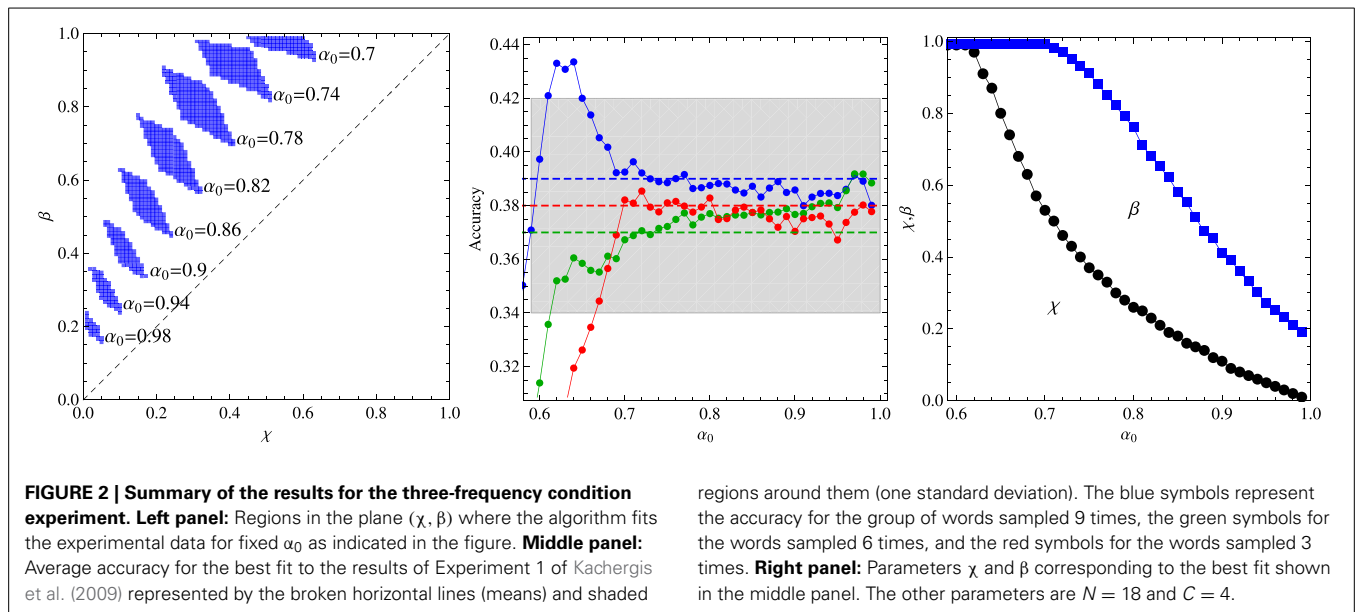
and the right panels exhibit the values of $\chi$ and $\beta$ corresponding to that fit. The broken horizontal lines and the shaded zones around them represent the means and standard deviations of the results of experiments carried out with 33 adult subjects (Kachergis et al., 2009).

It is interesting that although the words sampled more frequently are learned best in the two-frequency condition as expected, this advantage practically disappears in the three-frequency condition in which case all words are learned at equal levels within the experimental error. Note that the average accuracy for the words sampled 3 times is actually greater than the accuracy for the words sampled 6 times, but this inversion is not statistically significant, although, most surprisingly, the algorithm does reproduce it for $\alpha_0 \in [0.7, 0.8]$. According to Kachergis et al. (2009), the reason for the observed sampling frequency insensitivity might be because the high-frequency words are learned quickly and once they are learned subsequent trials containing those words will exhibit an effectively smaller within-trial ambiguity. In this vein, the inversion could be explained if by chance the words less frequently sampled were generally paired with the highly sampled words. Thus, contextual diversity seems to play a key role in cross-situational word learning.

## 4.2. CONTEXTUAL DIVERSITY AND WITHIN-TRIAL AMBIGUITY

In the first experiment aiming to probe the role of contextual diversity in the cross-situational learning, the 18 words were divided in two groups of 6 and 12 words each, and the contexts of size $C = 3$ were formed with words belonging to the same group only. Since the sampling frequency was fixed to 6 repetitions for each word, those words belonging to the more numerous group are exposed to a larger contextual diversity (i.e., the variety of different words with which a given word appear in the course of the training stage). The results summarized in **Figure 3** indicate clearly that contextual diversity enhances the learning accuracy. Perhaps more telling is the finding that incorrect responses are



**FIGURE 1 | Summary of the results for the two-frequency condition experiment. Left panel:** Regions in the plane ($\chi$, $\beta$) where the algorithm fits the experimental data for fixed $\alpha_0$ as indicated in the figure. **Middle panel:** Average accuracy for the best fit to the results of Experiment 1 of Kachergis et al. (2009) represented by the broken horizontal lines (means) and shaded regions around them (one standard deviation). The blue symbols represent the accuracy for the group of words sampled 9 times whereas the red symbols represent the accuracy for the words sampled 3 times. **Right panel:** Parameters $\chi$ and $\beta$ corresponding to the best fit shown in the middle panel. The other parameters are $N = 18$ and $C = 4$.

**FIGURE 2 | Summary of the results for the three-frequency condition experiment. Left panel:** Regions in the plane $(\chi, \beta)$ where the algorithm fits the experimental data for fixed $\alpha_0$ as indicated in the figure. **Middle panel:** Average accuracy for the best fit to the results of Experiment 1 of Kachergis et al. (2009) represented by the broken horizontal lines (means) and shaded regions around them (one standard deviation). The blue symbols represent the accuracy for the group of words sampled 9 times, the green symbols for the words sampled 6 times, and the red symbols for the words sampled 3 times. **Right panel:** Parameters $\chi$ and $\beta$ corresponding to the best fit shown in the middle panel. The other parameters are $N = 18$ and $C = 4$.



**FIGURE 3 | Summary of the results of the two-level contextual diversity experiment. Left panel:** Regions in the plane $(\chi, \beta)$ where the algorithm fits the experimental data for fixed $\alpha_0$ as indicated in the figure. **Middle panel:** Average accuracy for the best fit to the results of Experiment 2 of Kachergis et al. (2009) represented by the broken horizontal lines (means) and shaded regions around them (one standard deviation). The blue symbols represent the accuracy for the group of words belonging to the 12-components subgroup and the red symbols for the words belonging to the 6-components subgroup. All words are repeated exactly 6 times during the $t^* = 27$ learning trials. **Right panel:** Parameters $\chi$ and $\beta$ corresponding to the best fit shown in the middle panel. The other parameters are $N = 18$ and $C = 3$.

largely due to misassignments to referents whose words belong to the same group of the test word. In particular, Kachergis et al. (2009) found that this type of error accounts for 56% of incorrect answers when the test word belongs to the 6-components subgroup and for 76% when it belongs to the 12-components subgroup. The corresponding statistics for our algorithm with the optimal parameters set at $\alpha_0 = 0.9$ are 43% and 70%, respectively. The region in the space of parameters where the model can be said to describe the experimental data is greatly reduced in this experiment and even the best fit is barely within the error bars. It is interesting that, contrasting with the previous experiments, in this case the reinforcement procedure seems to play the more important role in the performance of the algorithm.

The effect of the context size or within-trial ambiguity is addressed by the experiment summarized in **Figure 4**, which is similar to the previous experiment, except that the words that compose the context are chosen uniformly from the entire repertoire of $N = 18$ words. Two context sizes are considered, namely, $C = 3$ and $C = 4$. In both cases, there is a large selection of parameter values that explain the experimental data, yielding results indistinguishable from the experimental average accuracies. This is the reason we do not exhibit a graph akin to those shown in the right panels of the previous figures. Since a perfect fitting can be obtained both for $\chi > \beta$ and for $\chi < \beta$, this experiment is uninformative with respect to these two abilities. As expected, increase of the within-trial ambiguity difficilitate

**FIGURE 4 | Summary of the results of the experiments where all words co-occur without constraint and the N = 18 words are repeated exactly 6 times during the t\* = 27 learning trials. Left panel:** Regions in the plane (χ, β) where the algorithm fits the experimental data for fixed $\alpha_0$ and context size C = 3. **Middle panel:** Same as the left panel but for context size C = 4. **Right panel:** Average accuracy for the best fitting of the results of Experiment 2 of Kachergis et al. (2009) represented by the broken horizontal lines (means) and shaded regions around them (one standard deviation). The red symbols are for C = 3 and the blue symbols for C = 4.

learning. In addition, the (experimental) results for C = 3 yield a learning accuracy value that is intermediary to those measured for the 6 and 12-components subgroups, which is in agreement with the conclusion that the increase of the contextual diversity enhances learning, since the mean number of different co-occurring words is 4.0 in the 6-components subgroup, 9.2 in the 12-components subgroup and 8.8 in the uniformly mixed situation (Kachergis et al., 2009).

### 4.3. FAST MAPPING

The experiments carried out by Kachergis et al. (2012) were designed to elicit participants' use of the mutual exclusivity principle (i.e., the assumption of one-to-one mappings between words and referents) and to test the flexibility of a learned word-object association when new evidence is provided in support to a many-to-many mapping. To see how mutual exclusivity implies fast mapping assume that a learner who knows the association $(w_1, o_1)$ is exposed to the context $\Omega = \{w_1, o_1, w_2, o_2\}$ in which the word $w_2$ (and its referent) appears for the first time. Then it is clear that a mutual-exclusivity-biased learner would infer the association $(w_2, o_2)$ in this single trial. However, a purely associative learner would give equal weights to $o_1$ and $o_2$ if asked about the referent of $w_2$.

In the specific experiment we address in this section, N = 12 words and their referents are split up into two groups of 6 words each, say $A = \{(w_1, o_1), \ldots, (w_6, o_6)\}$ and $B = \{(w_7, o_7), \ldots, (w_{12}, o_{12})\}$. The context size is set to C = 2 and the training stage is divided in two phases. In the early phase, only the words belonging to group A are presented and the duration of this phase is set such that each word is repeated 3, 6 or 9 times. In the late phase, the contexts consist of one word belonging to A and one belonging to B forming fixed couples, i.e., whenever $w_i$ appears in a context, $w_{i+6}$, with $i = 1, \ldots, 6$, must appear too. The duration of the late phase depends on the number of repetitions of each word that can be 3, 6, or 9 as in the early

phase (Kachergis et al., 2012). The combinations of the sampling frequencies yield 9 different training conditions but here we will consider only the case that the late phase comprises 6 repetitions of each word.
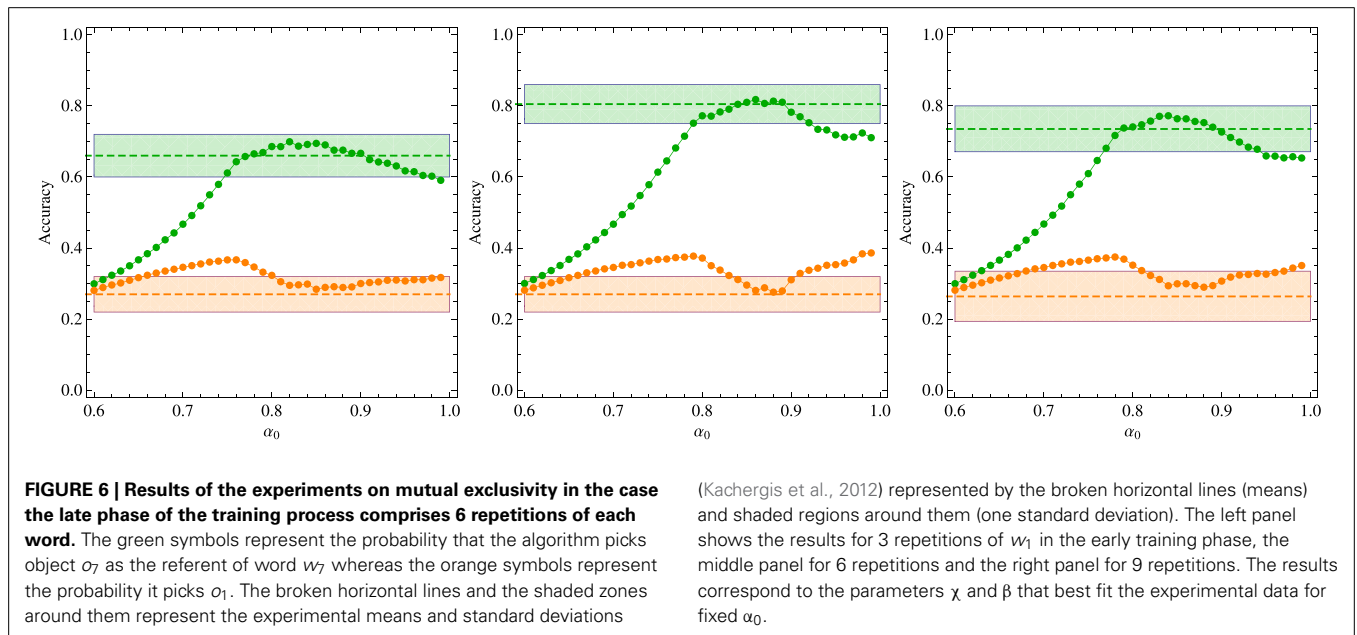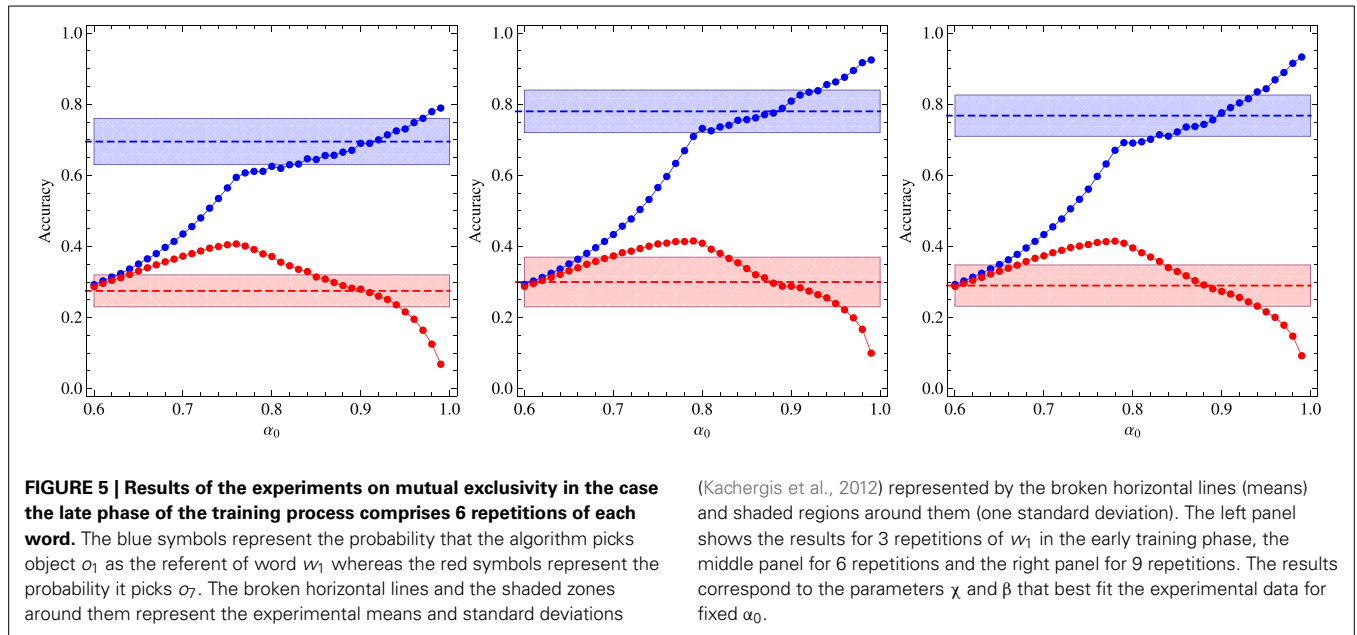
The testing stage comprises the play of a single word, say $w_1$, and the display of 11 of the 12 trained objects (Kachergis et al., 2012). Each word was tested twice with a time lag between the tests: once without its corresponding early object ($o_1$ in the case) and once without its late object ($o_7$ in the case). This procedure requires that we renormalize the confidences for each test. For instance, in the case $o_1$ is left out of the display, the renormalization is

$$P_{t*}{}'(w_1, o_j) = P_{t*}(w_1, o_j) / \sum_{o_k \neq o_1} P_{t*}(w_1, o_k) \qquad (15)$$

with $j = 2, \ldots, 12$ so that $\sum_{o_j \neq o_1} P_{t*}{}'(w_1, o_j) = 1$. Similarly, in the case $o_7$ is left out the renormalization becomes

$$P_{t*}{}'(w_1, o_j) = P_{t*}(w_1, o_j) / \sum_{o_k \neq o_7} P_{t*}(w_1, o_k) \qquad (16)$$

with $j = 1, \ldots, 6, 8, \ldots, 12$ so that $\sum_{o_j \neq o_7} P_{t*}{}'(w_1, o_j) = 1$. We are interested on the (renormalized) confidences $P_{t*}{}'(w_1, o_1)$, $P_{t*}{}'(w_1, o_7)$, $P_{t*}{}'(w_7, o_7)$, and $P_{t*}{}'(w_7, o_1)$, which are shown in **Figures 5, 6** for the conditions where words $w_i, i = 1, \ldots, 6$ are repeated 3 (left panel), 6 (middle panel), and 9 (right panel) times in the early learning phase, and the words $w_i, i = 1, \ldots, 12$ are repeated 6 times in the late phase. The figures exhibit the performance of the algorithm for the set of parameters χ and β that fits best the experimental data of Kachergis et al. (2012) for fixed $\alpha_0$. This optimum set is shown in **Figure 7** for the 6 early repetition condition, which is practically indistinguishable from the optima of the other two conditions. The conditions with the different word repetitions in the early phase intended to produce

**FIGURE 5 | Results of the experiments on mutual exclusivity in the case the late phase of the training process comprises 6 repetitions of each word.** The blue symbols represent the probability that the algorithm picks object $o_1$ as the referent of word $w_1$ whereas the red symbols represent the probability it picks $o_7$. The broken horizontal lines and the shaded zones around them represent the experimental means and standard deviations (Kachergis et al., 2012) represented by the broken horizontal lines (means) and shaded regions around them (one standard deviation). The left panel shows the results for 3 repetitions of $w_1$ in the early training phase, the middle panel for 6 repetitions and the right panel for 9 repetitions. The results correspond to the parameters $\chi$ and $\beta$ that best fit the experimental data for fixed $\alpha_0$.



**FIGURE 6 | Results of the experiments on mutual exclusivity in the case the late phase of the training process comprises 6 repetitions of each word.** The green symbols represent the probability that the algorithm picks object $o_7$ as the referent of word $w_7$ whereas the orange symbols represent the probability it picks $o_1$. The broken horizontal lines and the shaded zones around them represent the experimental means and standard deviations (Kachergis et al., 2012) represented by the broken horizontal lines (means) and shaded regions around them (one standard deviation). The left panel shows the results for 3 repetitions of $w_1$ in the early training phase, the middle panel for 6 repetitions and the right panel for 9 repetitions. The results correspond to the parameters $\chi$ and $\beta$ that best fit the experimental data for fixed $\alpha_0$.

distinct confidences on the learned association $(w_1, o_1)$ before the onset of the late phase in the training stage. The insensitivity of the results to these conditions probably indicates that association was already learned well enough with 3 repetitions only. Finally, we note that, though the testing stage focused on words $w_1$ and $w_7$ only, all word pairs $w_i$ and $w_{i+6}$ with $i = 1, \ldots, 6$ are strictly equivalent since they appear the same number of times during the training stage.

The experimental results exhibited in **Figure 6** offer indirect evidence that the participants have resorted to mutual exclusivity to produce their word-object mappings. In fact, from the perspective of a purely associative learner, word $w_7$ should be associated to objects $o_1$ or $o_7$ only, but since in the testing stage one of those objects was not displayed, such a learner would surely select the correct referent. However, the finding that $P_{t*}'(w_7, o_7)$ is considerably greater than $P_{t*}'(w_7, o_1)$ (they should be equal for an associative learner) indicates that there is a bias against the association $(w_7, o_1)$ which is motivated, perhaps, from the previous understanding that $o_1$ was the referent of word $w_1$. In fact, a most remarkable result revealed by **Figure 6** is that $P_{t*}'(w_7, o_7) < 1$. Since word $w_7$ appeared only in the late phase context $\Omega = \{w_1, o_1, w_7, o_7\}$ and object $o_1$ was not displayed in the testing stage, we must conclude that the participants produced spurious associations between words and objects that never appeared together in a context. Our algorithm accounts for these associations through Equation (4) in the case of new

words and, more importantly, through eqs. (9) and (13) due to the effect of the information efficiency factor $\alpha_t(w_i)$. The experimental data is well described only in the narrow range $\alpha_0 \in [0.85, 0.9]$.
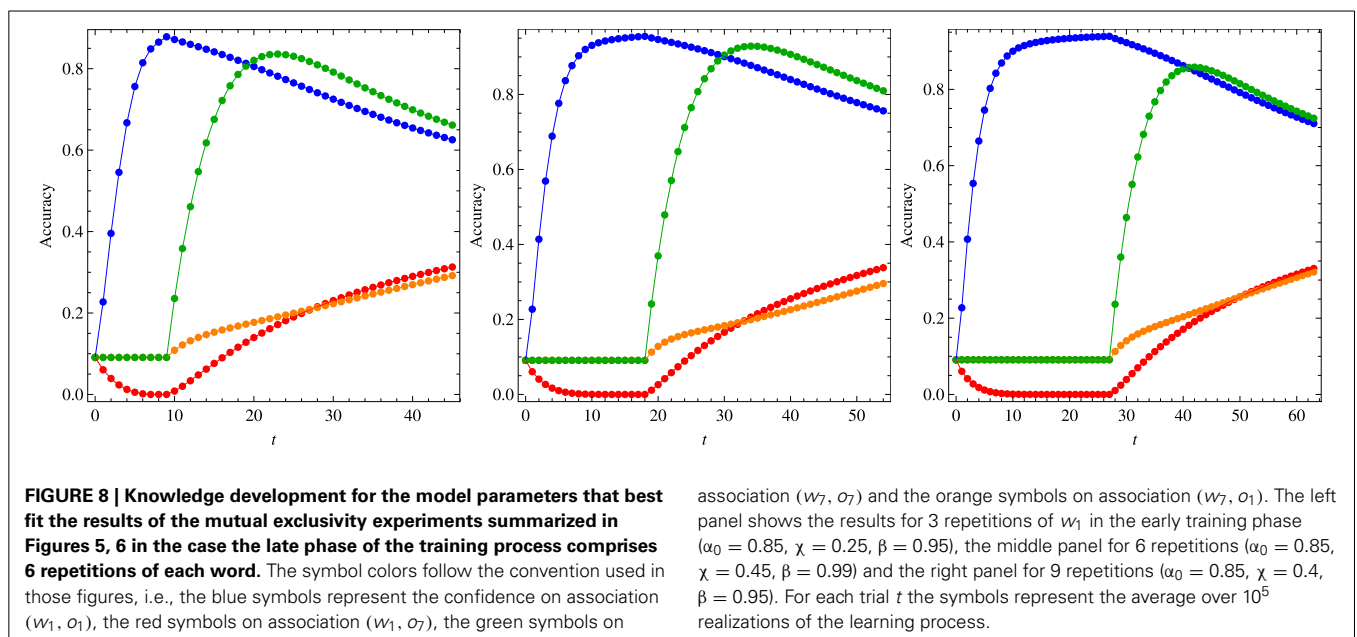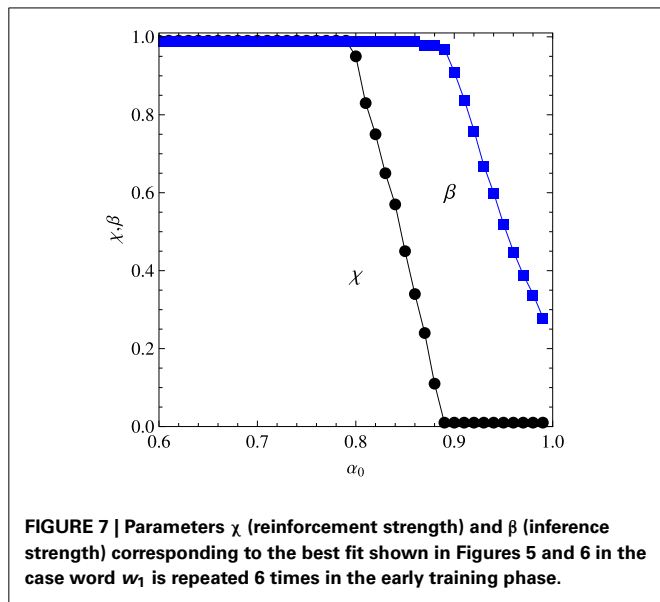
**Figure 8** exhibits the developmental timeline of the cross-situational learning history of the algorithm with the optimal set of parameters (see the figure caption) for the three different training conditions in the early training phase. This phase is characterized by the steady growth of the confidence on the association $(w_1, o_1)$ (blue symbols) accompanied by the decrease of the confidence on association $(w_1, o_7)$ (red symbols). As the word $w_7$ does not appear in the early training phase, the confidences on its association with any object remain constant corresponding to the accuracy value $1/11$ (we recall that $o_1$ is left out of the



**FIGURE 7 | Parameters $\chi$ (reinforcement strength) and $\beta$ (inference strength) corresponding to the best fit shown in Figures 5 and 6 in the case word $w_1$ is repeated 6 times in the early training phase.**

display in the testing stage). The beginning of the late training stage is marked by a steep increase of the confidence on the association $(w_7, o_7)$ (green symbols) whereas the confidence on $(w_1, o_1)$ decreases gradually. A similar gradual increase is observed on the confidence on the association $(w_7, o_1)$ (orange symbols). As expected, for large $t$ all confidences presented in this figure tend to the same value, since the words $w_1$ and $w_7$ always appear together in the context $\Omega = \{w_1, o_1, w_7, o_7\}$. Finally, we note that this developmental timeline is qualitatively similar to that produced by the algorithm proposed by Kachergis et al. (2012).

## 5. DISCUSSION

The chief purpose of this paper is to understand and model the mental processes used by human subjects to produce their word-object mappings in the controlled cross-situational word-learning scenarios devised by Yu and Smith (2007) and Kachergis et al. (2009, 2012). In other words, we seek to analyze the psychological phenomena involved in the production of those mappings. Accordingly, we assume that the completion of that task requires the existence of two cognitive abilities, namely, the associative capacity to create and reinforce associations between words and referents that co-occur in a context, and the non-associative capacity to infer word-object associations based on previous learning events, which accounts for the mutual exclusivity principle, among other things. In order to regulate the effectiveness of these two capacities we introduce the parameters $\chi \in [0, 1]$, which yields the reinforcement strength, and $\beta \in [0, 1]$, which determines the inference strength.

In addition, since the reinforcement and inference processes require storage, use and transmission of past and present information (coded mainly on the values of the confidences $P_t(w_i, o_j)$) we introduce a word-dependent quantity $\alpha_t(w_i) \in [0, 1]$ which gauges the impact of the confidences at trial $t - 1$ on the update of the confidences at trial $t$. In particular, the greater the certainty



**FIGURE 8 | Knowledge development for the model parameters that best fit the results of the mutual exclusivity experiments summarized in Figures 5, 6 in the case the late phase of the training process comprises 6 repetitions of each word.** The symbol colors follow the convention used in those figures, i.e., the blue symbols represent the confidence on association $(w_1, o_1)$, the red symbols on association $(w_1, o_7)$, the green symbols on association $(w_7, o_7)$ and the orange symbols on association $(w_7, o_1)$. The left panel shows the results for 3 repetitions of $w_1$ in the early training phase ($\alpha_0 = 0.85$, $\chi = 0.25$, $\beta = 0.95$), the middle panel for 6 repetitions ($\alpha_0 = 0.85$, $\chi = 0.45$, $\beta = 0.99$) and the right panel for 9 repetitions ($\alpha_0 = 0.85$, $\chi = 0.4$, $\beta = 0.95$). For each trial $t$ the symbols represent the average over $10^5$ realizations of the learning process.

about the referent of word $w_i$, the greater the relevance of the previous confidences. However, there is a baseline information gauge factor $\alpha_0 \in [0, 1]$ used to process words for which the uncertainty about their referents is maximum. The adaptive expression for $\alpha_t(w_i)$ given in Equation (7) seems to be critical for the fitting of the experimental data. In fact, our first choice was to use a constant information gauge factor (i.e., $\alpha_t(w_i) = \alpha \; \forall t, w_i$) with which we were able to describe only the experiments summarized in **Figures 1, 4** (data not shown). Note that a consequence of prescription (7) is that once the referent of a word is learned with maximum confidence (i.e., $P_t(w_i, o_j) = 1$ and $P_t(w_i, o_k) = 0$ for $o_k \neq o_j$) it is never forgotten.

The algorithm described in Section 3 comprises three free parameters $\chi$, $\beta$ and $\alpha_0$ which are adjusted so as to fit a representative selection of the experimental data presented in Kachergis et al. (2009, 2012). A robust result from all experiments is that the baseline information gauge factor is in the range $0.7 < \alpha_0 < 1$. Actually, the fast mapping experiments narrow this interval down to $0.85 < \alpha_0 < 0.9$. This is a welcome result because we do not have a clear-cut interpretation for $\alpha_0$—it encompasses storage, processing and transmission of information—and so the fact that this parameter does not vary much for wildly distinct experimental settings is evidence that, whatever its meaning, it is not relevant to explain the learning strategies used in the different experimental conditions. Fortunately, this is not the case for the two other parameters $\chi$ and $\beta$.

For instance, in the fast mapping experiments discussed in Subsection 4.3 the best fit of the experimental data is achieved for $\beta \approx 1$ indicating thus the extensive use of mutual exclusivity, and inference in general, by the participants of those experiments. Moreover, in that case the best fit corresponds to a low (but non-zero) value of $\chi$, which is expected since for contexts that exhibit two associations ($C = 2$) only, most of the disambiguations are likely to be achieved solely through inference. This contrasts with the experiments on variable word sampling frequencies discussed in Subsection 4.1, for which the best fit is obtained with intermediate values of $\beta$ and $\chi$ so the participants' use of reinforcement and inference was not too unbalanced. The contextual diversity experiment of Subsection 4.2, in which the words are segregated in two isolated groups of 12 and 6 components, offers another extreme learning situation, since the best fit corresponds to $\chi \approx 1$ and $\beta \approx 0$ in that case. To understand this result, first we recall that most of the participants' errors were due to misassignments of referents belonging to the same group of the test word, and those confidences were strengthened mainly by the reinforcement process. Second, in contrast to the inference process, which creates and strengthens spurious intergroup associations via Equation (12), the reinforcement process solely weakens those associations via Equation (9). Thus, considering the learning conditions of the contextual diversity experiment it is no surprise that reinforcement was the participants' choice strategy.

It is interesting to note that the optimal set of parameters that describe the fast mapping experiments (see **Figures 5–8**) indicate that there is a trade-off in the values of those parameters, in the sense that high values of the inference parameter $\beta$ require low values of the reinforcement parameter $\chi$. Since this is not an artifact of the model which poses no constrain on those values (e.g., they are both large for small $\alpha_0$), the trade-off may reveal a limitation on the amount of attentional resources available to the learner to distribute among the two distinct mental processes.

Our results agree with the findings of Smith et al. (2011) that participants use various learning strategies, which in our case are determined by the values of the parameters $\chi$ and $\beta$, depending on the specific conditions of the cross-situational word-learning experiment. In particular, in the case of low within-trial ambiguity those authors found that participants generally resorted to a rigorous eliminative approach to infer the correct word-object mapping. This is exactly the conclusion we reached in the analysis of the fast mapping experiment for which the within-trial ambiguity takes the lowest possible value ($C = 2$).

Although the adaptive learning algorithm presented in this paper reproduced the performance of adult participants in cross-situational word-learning experiments quite successfully, the deterministic nature of the algorithm hindered somewhat the psychological interpretation of the information gauge factor $\alpha_t(w_i)$. In fact, not only learning and behavior are best described as stochastic processes (Atkinson et al., 1965) but also the modeling of those processes requires (and facilitates) a precise interpretation of the model parameters, since they are introduced in the model as transition probabilities.

## REFERENCES

Atkinson, R. C., Bower, G. H., and Crothers, E. J. (1965). *An introduction to mathematical learning theory*. New York, NY: Wiley.

Baldwin, D. A., Markman, E. M., Bill, B., Desjardins, R. N., Irwin, J. M., and Tidball, G. (1996). Infants' reliance on a social criterion for establishing word-object relations. *Child Dev.* 67, 3135–3153. doi: 10.2307/1131771

Bloom, P. (2000). *How children learn the meaning of words*. Cambridge, MA: MIT Press.

Blythe, R. A., Smith, K., and Smith, A. D. M. (2010). Learning times for large lexicons through cross-situational learning. *Cogn. Sci.* 34, 620–642. doi: 10.1111/j.1551-6709.2009.01089.x

Cangelosi, A. (2010). Grounding language in action and perception: from cognitive agents to humanoid robots. *Phys. Life Rev.* 7, 139–151. doi: 10.1016/j.plrev.2010.02.001

Cangelosi, A., Tikhanoff, V., Fontanari, J. F., and Hourdakis, E. (2007). Integrating language and cognition: a cognitive robotics approach. *IEEE Comput. Intell. M.* 2, 65–70. doi: 10.1109/MCI.2007.385366

Fazly, A., Alishahi, A., and Stevenson, S. (2010). A probabilistic computational model of cross-situational word learning. *Cogn. Sci.* 34, 1017–1063. doi: 10.1111/j.1551-6709.2010.01104.x

Fontanari, J. F., and Cangelosi, A. (2011). Cross-situational and supervised learning in the emergence of communication. *Interact. Stud.* 12, 119–133. doi: 10.1075/is.12.1.05fon

Fontanari, J. F., Tikhanoff, V.. Cangelosi, A., Ilin, R., and Perlovsky, L. I. (2009). Cross-situational learning of object-word mapping using Neural Modeling Fields. *Neural Netw.* 22, 579–585. doi: 10.1016/j.neunet.2009.06.010

Frank, M. C., Goodman, N. D., and Tenenbaum, J. B. (2009). Using speakers' referential intentions to model early cross-situational word learning. *Psychol. Sci.* 20, 578–585. doi: 10.1111/j.1467-9280.2009.02335.x

Gleitman, L. (1990). The structural sources of verb meanings. *Lang. Acquis.* 1, 1–55. doi: 10.1207/s15327817la0101_2

Kachergis, G., Yu, C., and Shiffrin, R. M. (2009). "Frequency and contextual diversity effects in cross-situational word learning," in *Proceedings of 31st Annual Meeting of the Cognitive Science Society*. (Austin, TX: Cognitive Science Society), 755–760.

Kachergis, G., Yu, C., and Shiffrin, R. M. (2012). An associative model of adaptive inference for learning word-referent mappings. *Psychon. Bull. Rev.* 19, 317–324. doi: 10.3758/s13423-011-0194-6

Markman, E. M., and Wachtel, G. F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cogn. Psychol.* 20, 121–157. doi: 10.1016/0010-0285(88)90017-5

Newell, A. (1994). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

Pezzulo, G., Barsalou, L. W., Cangelosi, A., Fischer, M. H., McRae, K., and Spivey, M. J. (2013). Computational Grounded Cognition: a new alliance between grounded cognition and computational modeling. *Front. Psychol.* 3:612. doi: 10.3389/fpsyg.2012.00612

Pinker, S. (1990). *Language Learnability and Language Development*. Cambridge, MA: Harvard University Press.

Reisenauer, R., Smith, K., and Blythe, R. A. (2013). Stochastic dynamics of lexicon learning in an uncertain and nonuniform world. *Phys. Rev. Lett.* 110, 258701. doi: 10.1103/PhysRevLett.110.258701

Siskind, J. M. (1996). A computational study of cross-situational techniques for learning word-to-meaning mappings. *Cognition* 61, 39–91. doi: 10.1016/S0010-0277(96)00728-7

Smith, K., Smith, A. D. M., and Blythe, R. A. (2011). Cross-situational learning: an experimental study of word-learning mechanisms. *Cogn. Sci.* 35, 480–498. doi: 10.1111/j.1551-6709.2010.01158.x

Smith, L. B., and Yu, C. (2013). Visual attention is not enough: individual differences in statistical word-referent learning in infants *Lang. Learn. Dev.* 9, 25–49. doi: 10.1080/15475441.2012.707104

Steels, L. (2003). Evolving grounded communication for robots. *Trends Cogn. Sci.* 7, 308–312. doi: 10.1016/S1364-6613(03)00129-3

Tilles, P. F. C., and Fontanari, J. F. (2012a). Minimal model of associative learning for cross-situational lexicon acquisition. *J. Math. Psychol.* 56, 396–403. doi: 10.1016/j.jmp.2012.11.002

Tilles, P. F. C., and Fontanari, J. F. (2012b). Critical behavior in a cross-situational lexicon learning scenario. *Europhys. Lett.* 99, 60001. doi: 10.1209/0295-5075/99/60001

Vogt, P. (2012). Exploring the robustness of cross-situational learning under Zipfian distributions. *Cogn. Sci.* 36, 726–739. doi: 10.1111/j.1551-6709.2011.1226.x

Waxman, S. R., and Gelman, S. A. (2009). Early word-learning entails reference, not merely associations. *Trends Cogn. Sci.* 13, 258–263. doi: 10.1016/j.tics.2009.03.006

Yu, C., and Smith, L. B. (2007). Rapid word learning under uncertainty via cross-situational statistics. *Psychol. Sci.* 18, 414–420. doi: 10.1111/j.1467-9280.2007.01915.x

Yu, C., and Smith, L. B. (2008). "Statistical cross-situational learning in adults and infants," in *Proceedings of the 32th Annual Boston University Conference on Language Development* (Somerville, MA: Cascadilla Press), 562–573.

Yu, C., and Smith, L. B. (2012a). Modeling cross-situational word-referent learning: prior questions. *Psychol. Rev.* 119, 21–39. doi: 10.1037/a0026182

Yu, C., and Smith, L. B. (2012b). Embodied attention and word learning by toddlers. *Cognition* 125, 244–262. doi: 10.1016/j.cognition.2012.06.016