



HAL
open science

Familiar words can serve as a semantic seed for syntactic bootstrapping

Mireille Babineau, Alex A de Carvalho, John Trueswell, Anne Christophe

► To cite this version:

Mireille Babineau, Alex A de Carvalho, John Trueswell, Anne Christophe. Familiar words can serve as a semantic seed for syntactic bootstrapping. *Developmental Science*, 2020, 24 (1), pp.e13010. 10.1111/desc.13010 . hal-02951541

HAL Id: hal-02951541

<https://hal.science/hal-02951541>

Submitted on 10 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Familiar words can serve as a semantic seed for syntactic bootstrapping

Mireille Babineau a,b *, Alex de Carvalho c, John Trueswell d, and Anne Christophe a,b

a. Laboratoire de Sciences Cognitives et Psycholinguistique, DEC-ENS / EHESS / CNRS,

Ecole Normale Supérieure – PSL University, Paris, France

b. Maternité Port-Royal, AP-HP, Faculté de Médecine Paris Descartes, France

c. Université de Paris, LaPsyDÉ, CNRS F-75005 Paris, France

d. Department of Psychology, University of Pennsylvania, Philadelphia, PA, United States of
America

Corresponding author: correspondence concerning this article should be addressed to Mireille Babineau. Email: mireillebbineau@gmail.com; Address: Laboratoire de Sciences Cognitives et Psycholinguistique, École normale supérieure, 29 rue d'Ulm, 75005 Paris, France.

Research highlights

- We ask how children come to do syntactic bootstrapping, i.e. how they managed to learn which features are shared by words occurring in specific syntactic contexts.
- In our study, children learned the distribution of a novel function word (“ko” before nouns or before verbs) after watching a 5-min training video.
- Children inferred the probable meaning of novel content words co-occurring with the newly-learned function word, interpreting them as either referring to novel actions or objects.
- Young children can go about identifying semantic properties associated with specific syntactic contexts from tracking the occurrence of a small set of familiar words.

Abstract

Young children can exploit the syntactic context of a novel word to narrow down its probable meaning. But how do they learn which contexts are linked to which semantic features in the first place? We investigate if 3-to-4 year old children (n=60) can learn about a syntactic context from tracking its use with only a few familiar words. After watching a 5-min training video in which a novel function word (i.e. “ko”) replaced either personal pronouns or articles, children were able to infer semantic properties for novel words co-occurring with the newly-learned function word (i.e. objects vs actions). These findings implicate a mechanism by which a distributional analysis, associated with a small vocabulary of known words, could be sufficient to identify some properties associated with specific syntactic contexts.

Keywords: Word learning, Language development, Language processing, Syntactic bootstrapping, Eye movements

Familiar words can serve as a semantic seed for syntactic bootstrapping

Research on language development has provided abundant evidence on how children learn the meaning of new words by relying on the other words in the sentence (the syntactic context), a mechanism called *syntactic bootstrapping* (e.g., Arunachalam, 2016; Bernal, Lidz, Millotte, & Christophe, 2007; Fisher, Gertner, Scott, & Yuan, 2010; Gillette, Gleitman, Gleitman, & Lederer, 1999; Gleitman, 1990; Landau & Gleitman, 1985). For instance, when hearing “It’s a pratch!” 18-month-olds can infer that *pratch* refers to an object, while upon hearing “It’s pratching!,” they can infer that it refers to an action (de Carvalho, He, Lidz, & Christophe, 2019; He & Lidz, 2017). Syntactic contexts are often revealed by functional elements, such as function words and morphemes (e.g., “the”, “a”, “she”, “he”, “-ing”), which are so frequent that infants store and recognize them before the age of one (e.g., in English: Shi et al., 2006; in French: Shi & Lepage, 2008). Crucially, function words can be exploited during syntactic and semantic acquisition, since they provide useful information about the co-occurring content words, first enabling infants to assign words that they do not know yet to grammatical categories (e.g., Babineau, Shi & Christophe, under review; Höhle, Weissenborn, Kiefer, Schulz, & Schmitz, 2004; Shi & Melançon, 2010) and then to infer their probable meanings (e.g., Bernal et al., 2007; de Carvalho et al., 2019a; He & Lidz, 2017). That is, young children expect a novel word following determiners to be a noun (i.e. mapping it to an object), and a novel word following pronouns to be a verb (i.e. mapping it to an action). It is clear that children can make use of the syntactic contexts that they already know to speed up their acquisition of vocabulary very early in their development, but it is less clear how they have gotten to this point.

The question of how children learned which syntactic contexts correspond to which semantic features is still unanswered. Recent findings about infants’ abilities can give us an insight on the learning mechanism at play. On the one hand, infants are known to have the

ability to track statistical regularities, which should help them to learn some aspects of the grammar of their native language, as they can do after a brief exposure to an artificial language (e.g. between the age of 12-18 months: Gerken, Wilson, & Lewis, 2005; Gómez & Gerken, 1999). Distributional patterns are present in natural languages, i.e. specific functional elements tend to co-occur with content words from specific word classes. For instance, articles tend to precede nouns while personal pronouns tend to precede verbs in languages such as English and French. Attempts at modeling the distributional learning of syntactic categories via the co-occurrences of content words with specific function words have consequently been successful (e.g., Chemla, Mintz, Bernal, & Christophe, 2009; Mintz, Newport, & Bever, 2002; Mintz, 2003; Weisleder & Waxman, 2010). On the other hand, rudimentary knowledge of the meaning of words develops early during infancy: preverbal infants have managed to learn something about the meaning of frequent and concrete words (e.g. “banana”, “feet”; Bergelson & Swingley, 2012; Tincoff & Jusczyk, 1999), possibly through their multimodal daily experience, in addition to noticing the semantic relatedness between those words (Bergelson & Aslin, 2017). By the end of their first year, they even group concepts into basic semantic categories (e.g., object, action, agent; Carey, 2009). Given these results in the domain of early syntactic and semantic acquisition, and the possible synergies between these two during language development, researchers that have modeled distributional learning of categories have proposed that children may be using their knowledge of a handful of word meanings as a *seed* for future syntactic categories (e.g., Christophe, Dautriche, de Carvalho, & Brusini, 2016; Gutman, Dautriche, Crabbé, & Christophe, 2014; see also: Christodoulopoulos, Roth, & Fisher, 2016; Connor, Fisher & Roth, 2012; Connor, Gertner, Fisher & Roth, 2010; Fisher & Gleitman, 2002; Gleitman, Cassidy, Nappa, Papafragou, & Trueswell, 2005). Based on the *semantic seed hypothesis*, young children would track the syntactic contexts in which known familiar words occur through a distributional analysis (e.g., “**the** car”; “**the** ball”), then use this information to

infer some semantic properties of novel words that they encounter in the same syntactic contexts: “**the** dax” → “dax” shares some characteristics with “car” and “ball” (and possibly refers to an object). This mechanism would enable the creation of syntactic categories, which in turn would guide children’s interpretation of novel word meanings.

The present study tested for the first time whether young children can indeed learn about syntactic categories by paying attention to the contexts in which familiar words appear. To do so, we taught children a novel function word embedded in sentences containing known content words (the *seed*). A new function word, “ko”, was introduced as a novel syntactic context during a short familiarization phase (i.e. a 5-minute training video). The goal was to assess whether children can use their ability to compute distributional regularities to track the co-occurrence of this new function word with a specific type of familiar words (either nouns or verbs, in a between-participants design), and then use this information when encountering novel words co-occurring with the newly-learned function word. To test this, we relied on an experimental design recently developed by de Carvalho, Babineau, Trueswell, Waxman, & Christophe (2019), who showed that 3-to-4-year-olds can use known function words in real-time to predict the syntactic category of novel content words. In this task participants saw two videos side-by-side on a TV-screen: one showing a person performing a novel action, and the other a person passively holding a novel object. At the same time, participants heard a novel word preceded either by a determiner (e.g., *Regarde! Une dase!* – “Look! A dase!”) or a pronoun (e.g., *Regarde! Elle dase!* – “Look! She’s dasing!”). Three-to-4-year-olds exploited function words online to categorize novel words and infer their meanings: they looked more to the novel action in the Verb condition, while participants in the Noun condition looked more to the novel object.

In the current experiment, we inquired whether 3-4 year-olds could use what they learned of the novel function word “ko” to narrow down the probable meaning of a novel

content word, e.g. *dase* in *ko dase*, referring either to a novel object (if “ko” was heard in a noun context) or a novel action (if “ko” was heard in a verb context). We hypothesized that children’s training condition would influence their implicit visual attention during the test trials. That is, children in the Verb condition would learn to attend to actions after hearing the novel word (e.g. “ko”, such as while children in the Noun condition would attend to objects. Hence, when listening to a test sentence in which a novel content word (e.g. *dase*) is preceded by the novel function word (e.g., *Regarde! Ko dase!* “Look! Ko dase”), participants who heard “ko” before verbs (Verb condition) should look more to the novel action, while participants who heard “ko” before nouns (Noun condition) should look more to the novel object. Moreover, we recorded participants’ pointing responses, as a measure of their explicit guesses regarding the meaning of the novel words. We also tested adult university students as a control group.



Training video

<u>Noun condition:</u>	<u>Verb condition:</u>
<p><i>Ko cochon, il veut encore jouer avec ko jolie tortue! Il fait rouler ko balle comme ça.</i> “Ko pig, he wants to play again with ko pretty turtle! He rolls ko ball this way”.</p>	<p><i>Le cochon, ko veut encore jouer avec la jolie tortue! Ko fait rouler la balle comme ça.</i> “The pig, ko wants to play again with the pretty turtle! Ko rolls the ball this way”.</p>

Figure 1: Example of the script used in the two versions of the training video

Method

The method, analyses and criteria for exclusion of participants were pre-registered on the OSF (Open Science Framework) database before running the experiment (the formal preregistration can be accessed with the following link:

https://osf.io/6s9eh/?view_only=c65b329b9b2548309114548c2b87b8d7). The materials, collected data, and data analyses are freely available to readers through the following link:

https://osf.io/79j53/?view_only=0197429abd4e4bd5b451bd127911f642.

Participants

Preschoolers. Sixty 3- to 4-year-old monolingual French-speaking children (3;3 to 4;4, M age = 3;8, 33 females) participated to the study (30 in each condition). A total of 25 children were tested at the laboratory and the remaining 35 were tested in two public preschools in Paris. Parents signed an informed consent form. The study was approved by the local ethics committee (CER Paris Descartes), as well as by the school boards.

An additional 27 children were tested, but not included in the final analysis because they failed to comply during the experiment, e.g. fussing, not pointing or taking off the headphones (17), because they were bilingual (6), or because of technical problems (4). As stated in the pre-registration, we removed from the statistical analyses test trials with more than 25% missing data frames. If a participant had more than two excluded trials (out of 4), she was excluded from eye-tracking data analysis (eight children). Four additional participants had corrupted eye-tracking files. Hence, a total of 12 children were excluded from the eye-tracking analyses, resulting in 48 children remaining for those analyses. The preregistered number of participants sufficient for the eye-tracking analyses was based on de Carvalho et al. (2019), which found a medium-size effect with a final sample of 48 children. Note that in our study, pointing responses were kept for subjects with low-quality/corrupted eye-tracking data. In fact,

in order to complete the experiment and to be included in the final sample, children had to be engaged in the task and point toward one of the still videos on every trial when the experimenter prompted them to do so. Hence, there are no pointing data missing in the final sample of 60 children.

Adults. Forty-nine native French-speaking adults (18 to 34 years of age, M age = 23; 28 females) participated in the control study (25 in the noun condition, 24 in the verb condition). They were tested in the laboratory, where they signed an informed consent form.

Five additional participants were tested, but not included in the final analysis. Three were excluded based on their poor eye-tracking data (see the description of the criteria above), and two participants were not included in the final analysis because they had seen the novel animals or heard about the training videos before the experiment.

Stimuli and design

In a between-subject design, participants were assigned to one of two conditions (Noun vs Verb condition). In the first part of the experiment, they watched a training video showing a female native French speaker (last author) using child-friendly speech as she acted out scenarios with stuffed animals and toys. We created two versions of the training video (see Figure 1). In both videos, the same script was used with the exception that articles *un/une/le/la* (feminine and masculine forms of “a”/“the”) were replaced by “ko” in the video assigned to the Noun Condition (“e.g., “ko pretty turtle”, “ko pig”), while pronouns *elle/il* (“she”/ “he”) were replaced by “ko” in the video assigned to the Verb Condition (e.g. “ko will play”, “ko rolls”; see Figure 1). The new function word “ko” was presented 60 times in total in each of the videos, co-occurring with 14 different familiar nouns in the Noun condition (e.g. *biberon* “bottle”, *chat* “cat”, *cochon* “pig”, *livre* “book”) and with 17 different familiar verbs in the Verb condition (e.g. *donner* “to give”, *jouer* “to play”, *regarder* “to look”, *tomber* “to fall”).

To facilitate the segmentation of the new function word and the recognition of the familiar co-occurring words, syntactic contexts varied for both noun and verb contexts. Specifically, to avoid perceiving the Article + N sequence as a whole new word (e.g., *kochat* “kocat” as a word), both Article+N and Article+Adj+N structures were used in the Noun condition (e.g. *ko joli chat* ‘ko pretty cat’) for 10 of the 14 familiar nouns. For the same reason, different verb tenses were used for 14 of the 17 familiar verbs, with 11 verbs appearing in two verb tenses (either present and past or present and future), and 3 appearing in the present, future, and past tense (e.g. *danser* ‘to dance’ appeared in *ko va danser* ‘ko will dance’, *ko danse* ‘ko dances’, *ko a dansé* ‘ko danced’). The full script of the videos along with summary tables are available on the OSF page, as well as the entire training videos. In both conditions, the novel function word “ko” shared similar acoustic properties as the replaced real function words. For instance, although “ko” tended to be shorter in duration than real function words (both articles and pronouns), it was not significantly so ($p > .1$). Note that the short duration of “ko” might be due to the fact that the closure before [k] was not counted as part of the word since it was not possible to distinguish it from a pause. The detailed results of the speech analysis done on half of the training videos can be found in the additional material section of the OSF.

In the second part of the experiment (i.e., the test phase), participants were tested with two videos displayed side-by-side on a TV-screen: one video showing an agent performing an intransitive action, and the other video showing an agent passively holding an object (the same videos used in de Carvalho et al., 2019, but with new sound tracks recorded by the last author). A total of six trials were presented to each participant: two practice trials followed by four test trials. The practice trials included two pairs of videos referring to familiar words (*une voiture* “a car” vs *elle dort* “she is sleeping”; *un ballon* “a ball” vs *il mange* “he is eating”). Every participant had one practice trial in which the target was a familiar verb and one in which the target was a familiar noun. The new function word “ko” replaced the article in the Noun

condition and the pronoun in the Verb condition (e.g. a participant in the Noun condition might hear *ko voiture* “ko car” for the first practice trial and *il mange* “he is eating” for the second one; while a participant in the Verb condition might hear *ko dort* “ko is sleeping” for the first practice trial and *un ballon* “a ball” for the second one). The practice trials not only aimed at familiarizing participants with the structure of the test phase (e.g. two videos would be presented side-by-side and only one of them would be the target), but also to extend the teaching period of the novel function word by presenting simple visual scenes with restricted mappings. The test trials included four pairs of videos referring to novel words (i.e., *dase*, *nuve*, *rane*, *fome*). Note that a third-person singular verb is unmarked in French (e.g., *elle marche* “she is walking”), as is a singular noun (e.g., *une marche* “a step”) making it possible to create ambiguous novel words that could be either a noun or a verb.

All trials followed the same structure (see Figure 2), aimed at familiarizing participants with the content of the videos, and to reduce their novelty and potential biases. Each of the two videos were presented alone on each side of a TV-screen during an inspection period. Note that the practice trials contained an informative audio prompt such as “*Oh regarde! Il mange. Tu as vu ça?*” - “*Oh look! He’s eating. Did you see that?*” during the inspection period, whereas the test trials contained uninformative audio prompts such as “*Oh regarde! Tu as vu ça?*” - “*Oh look! Did you see that?*”. Then, during the contrast period, both videos were presented side-by-side during 10 sec along with an uninformative audio prompt (e.g. *Et voilà les deux!* “And here’s the two of them”). Subsequently, a fixation point was presented during 6 sec along with an informative audio prompt in which the target was named once (e.g., *Oh regarde! Ko dase!* “Oh look! Ko dase!”). The two videos finally reappeared side-by-side during 10 sec together with two more repetitions of the target (e.g., *Tu vois? Ko dase! Wow regarde! Ko dase.* “Do you see that? Ko dase! Wow look! Ko dase.”) meant to attract the participant’s attention toward one of the videos. The two videos would freeze after 10 sec, indicating to the

participant that she needed to point towards the video that the woman referred to. After each trial, a picture of a baby (along with a sound of laughter) would appear in the center of the screen (5 sec). Note that the test phase with novel content words presented the exact same stimuli (videos and sentences) to all participants, in both conditions (Noun and Verb). Hence, the only difference between these two conditions was the training phase and the practice trials.

Apparatus and Procedure

Participants were tested individually, with children being tested either in a sound-attenuated booth in the laboratory (with their caregiver sitting behind them) or in a quiet room at their preschool, and adults in the control group being tested at the laboratory. Participants sat in front of a 27-in TV-screen positioned approximately 70 cm away from them. Their eye-gaze towards the videos was recorded by an EYELINK-1000, with a time-sample collected every 2 ms. A 5-point calibration procedure was used.

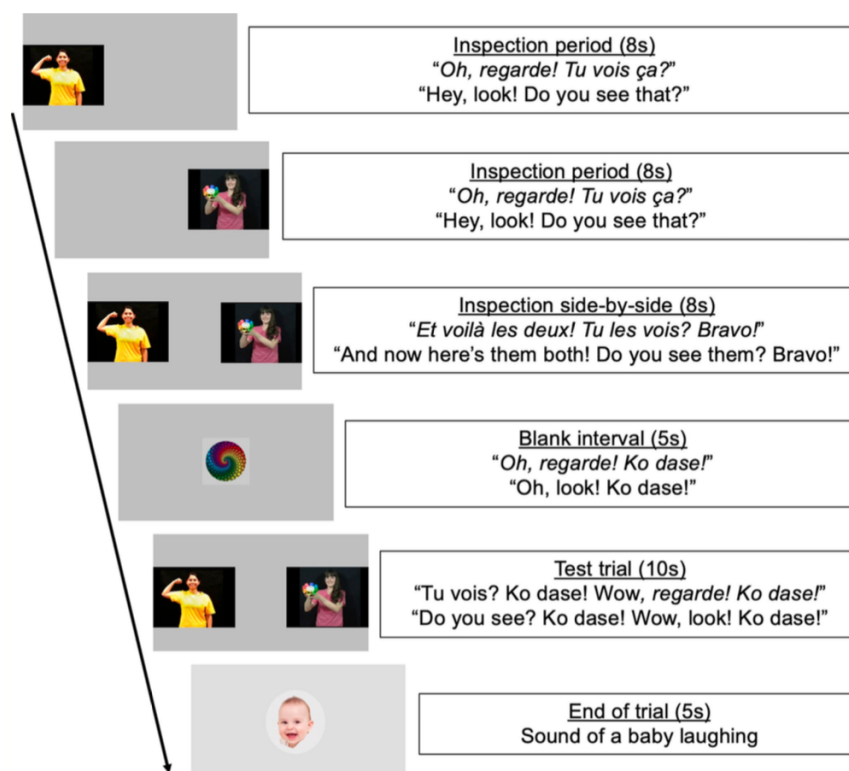


Figure 2: Example of the time-course of a test trial. The practice trials were presented in the same way with the exception that during the inspection period the prompt sentences were informative: they contained the familiar words (e.g., 'Oh look! Ko ball! Do you see that?'). Children were asked to point after the test trial has ended, while the 'frozen' version of the videos were displayed on the screen. Once their response was entered, the end-of-trial video with the baby laughing was presented

Children were given verbal instructions before starting the task. They were told that they would need to wear headphones so that they could hear a woman tell them a story. After the stories, they would watch videos and the woman would tell them where to look. Since the experimenter would not have headphones, she would not know which video the woman talked about. Hence, children were asked to point to the correct video (left vs right) to let the experimenter know what the woman talked about. The experimenter reminded children to point during the first practice trial. She gave no feedback about whether their answer was correct or not (and she did not herself know which answer was correct). Adults were told that they were participating in an experiment intended for children, and they received the same instructions as the children. The first fourteen adults tested were not asked to point after the end of each trial (only their eye-gaze was recorded). For the remaining participants, the experimenter gathered pointing responses after each trial, just as for children. At the end of the experiment, adult participants answered questions regarding their degree of awareness and knowledge of the novel function word's usage.

The experiment was composed of two phases: a familiarization phase in which a training video was presented (either the version for the Noun condition or the one for Verb condition), and a test phase including two practice trials with familiar words, and four test trials with novel words. During the test phase, participants needed to point toward the target video in order to proceed to the next trial. The experimenter recorded participants' choice by pressing a key on the keyboard of the computer.

Data processing and analyses

We report our three pre-registered analyses looking at (1) the time-course of participants' eye-gaze, reflecting their real-time interpretation of test sentences, (2) the looking times averaged over the entire duration of the test trials, reflecting their overall interpretation

of test sentences and (3) the pointing responses, reflecting participants' final interpretation of the target novel words. Additional analyses which were not preregistered or which look at various subsets of the data (such as practice trials) can be found on OSF following this link: https://osf.io/79j53/?view_only=0197429abd4e4bd5b451bd127911f642. We used the package `eyetrackingR` (Dink & Ferguson, 2016) to conduct the analyses, and the `ggplot2` package (Wickham, 2009) to plot the eye-gaze data and the pointing responses. We down-sampled the eye-tracking data (initially collected every 2 ms) by averaging to one sample every 20 ms. We removed from the statistical analyses test trials with more than 25% missing data frames, i.e. 30 trials for preschoolers (out of 192) and 27 trials for adult participants (out of 196).

To test the effect of Condition (Noun vs Verb), we ran a cluster-based permutation analysis for each group (as in de Carvalho, Dautriche, & Christophe, 2017; Hahn, Snedeker, & Rabagliati, 2015; Von Holzen & Mani, 2012; see Maris & Oostenveld, 2007 for a formal presentation of this analysis), to find whether there were time-windows (i.e. clusters) during which the two conditions were significantly different from each other. This conservative analysis has the advantage of avoiding fixing a time-window arbitrarily beforehand. Since the fixations to the two videos were complementary (apart from the time spent looking away, which is not significantly different between conditions), the analysis used the proportion of fixations toward the action video as the dependent variable. The analysis was conducted on the entire test trial (0-10s), since participants heard the target words once before the beginning of the test trials (during the period with the fixation point), which could already have an impact on their eye-gaze. The steps of the analysis are as follows. For each time point, a *t*-test testing for the effect of Condition (Noun vs Verb condition) was conducted on the proportion of looks toward the action video. Adjacent time points that have a *t*-value greater than our predefined threshold ($t = 1.5$) form a cluster, and its size is the sum of the *t* values at each time point within this cluster. Crucially, to test the probability of observing a cluster of that size by chance, we

conducted 1000 simulations where we randomly shuffle the conditions (Noun, Verb). For each simulation, the same procedure mentioned above was used to calculate the size of the biggest cluster. Finally, the cluster found with our real data is considered significant if its size is greater than the size of the largest cluster found in 95% of the simulations (ensuring a p-value of .05).

Results

Results of the eye-tracking data

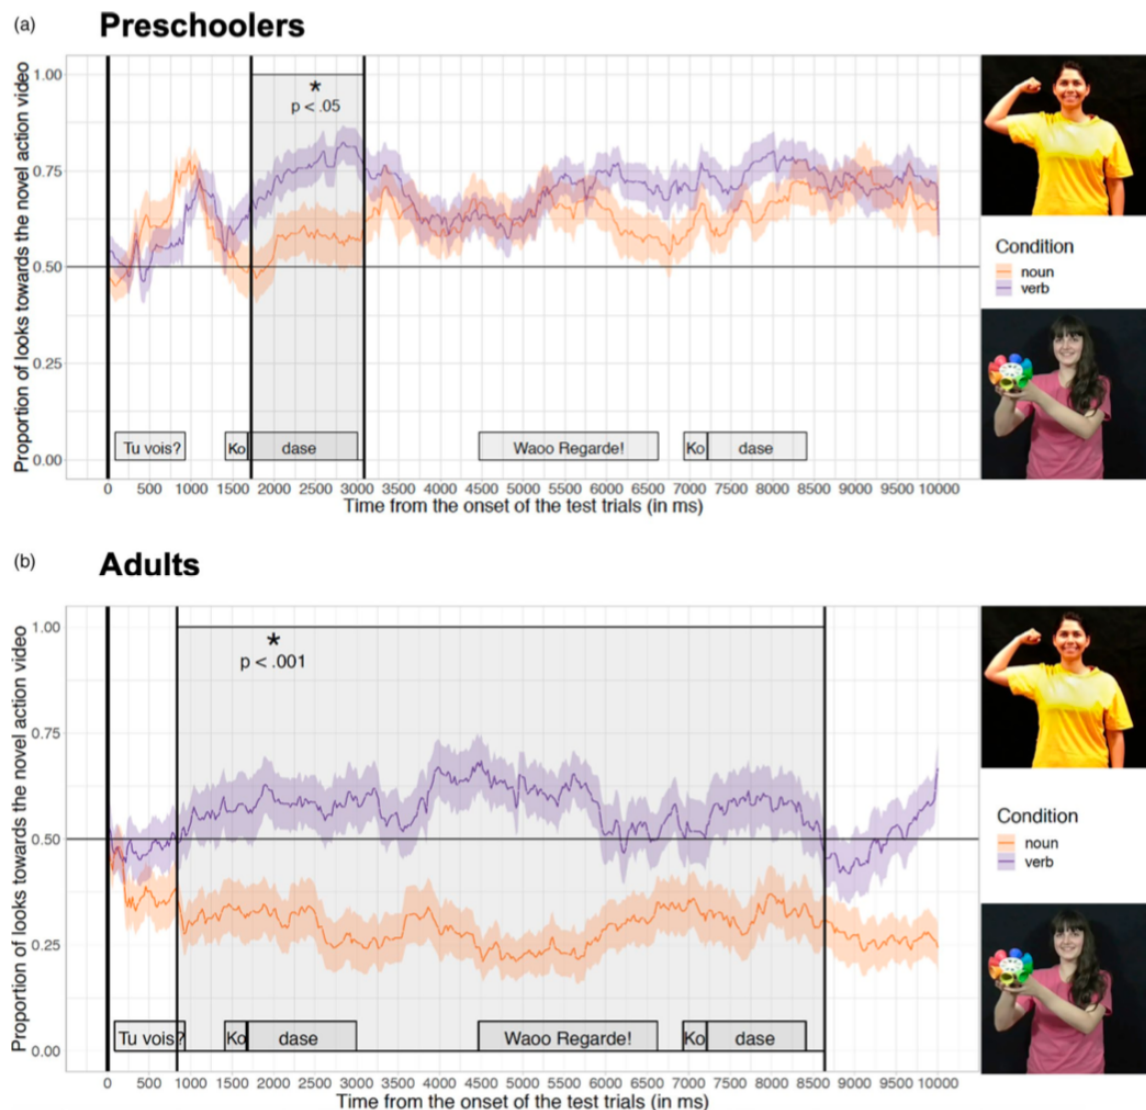


Figure 3: Proportion of looks toward the action video for preschoolers (a) and adults (b), time-locked to the onset of the trial for participants in the Verb condition (purple curve) and those in the Noun condition (orange curve). Error bars represent the standard error of the mean. A cluster-based permutation test revealed significant difference between the Verb and the Noun conditions for both age groups (grey time-windows). For preschoolers, the significant time-window started 307 ms after the first occurrence of the novel function word ‘ko’. For adults, the significant time-window started slightly after the onset of the trial. Note that a first sentence containing the novel word preceded by ‘ko’ was played during the blank interval.

Figure 3 shows children's average proportion of looks toward the action video in the Verb condition (purple curve) and in the Noun condition (orange curve), time-locked to the beginning of trial onset. Visual inspection of the data reveals that overall children tended to look more toward the action videos, most probably due to the movements in these videos, which might have attracted their attention. In the action videos (illustrating the verb interpretation), the actors were moving their body and/or arms to perform a novel action, whereas in the object videos (illustrating the noun interpretation), the actors were simply holding novel objects and looking at them. Crucially, after processing the first occurrence of “ko” during the test trial (and second occurrence overall), children in the Verb condition increased their looks toward the action video more than children in the Noun condition. Aligning with these observations, the cluster-based analysis found a significant time-window where the proportion of looks toward the action video was significantly different in the Verb condition compared to the Noun condition, from 1720ms after the beginning of the trial until 3080ms ($p = .043$).

Figure 4 shows adults' average proportion of looks toward the action video in the Verb condition (purple curve) and in the Noun condition (orange curve), time-locked to the beginning of trial onset. Visual inspection of the data reveals that from the beginning of the test trials participants in the Verb condition increased their looks toward the action video more than those in the Noun condition. Note that a first occurrence of “ko” with a novel word was presented during the blank interval before the onset of each test trial, enabling adults to anticipate the side of the target video. Aligning with these observations, the cluster-based permutation analysis found a significant time-window where the proportion of looks toward the action video was significantly different between the Verb and the Noun conditions, from 840ms after the beginning of the trial until 8640ms ($p < .001$). This shows that adults, just

like 3-to-4-year-olds, were able to exploit the novel function word in real-time to infer the probable meaning of the novel co-occurring content words.

We also compared looking times averaged over the entire duration of the test trials (see children's data in Figure 5, and adults' data in Figure 6). This analysis aimed to ensure that an effect would not be missed simply because there is no single moment at which all participants direct their gaze to one of the videos. To do so, a two sample t-test was conducted on the average overall looking time per participants. Children from the Verb Condition looked significantly more toward the action video ($M=.691$, $SD=.095$) than children from the Noun Condition ($M=.626$, $SD=.0858$; $t(46)=-2.48$, $p=.017$; Cohen's $d = .716$). Similarly, adults from the Verb Condition looked significantly more toward the action video ($M=.559$, $SD=.216$) than those from the Noun Condition ($M=.297$, $SD=.22$; $t(46)=-4.21$, $p=.001$; Cohen's $d = -1.202$). These results confirm that both children and adults from the Verb condition were more likely to interpret the novel words as referring to novel actions than those in the Noun condition.

Results of the pointing responses

The proportion of pointing towards the action video for each child is shown in Figure 7, and for each adult in Figure 8. To analyze this data, we ran for each group a mixed model analysis in which we modeled the occurrence of a pointing response toward the action video as predicted by the condition (Noun Condition coded as 0, Verb Condition coded as 1). With our model, which included a by-subject intercept, we found a significant main effect of Condition ($\beta = .92$; $SE=.34$; $z = 2.71$; $p = .007$), predicting an increase of 0.22 in the probability of children pointing to the action video if they were in the Verb Condition. As can be seen in the figure 7, children in the Verb Condition pointed more toward the action video ($M=.55$; $SEM=.05$) than those in the Noun Condition ($M=.35$; $SEM: .04$). This trend was found for all of the four novel target words. With the adult data, we also found a significant main effect of

Condition ($\beta = 2.39$; $SE = .62$; $z = 3.84$; $p < 0.001$), predicting an increase of 0.53 in the probability of adult participants pointing to the action video if they were in the Verb Condition.

As can be seen in figure 8, adults in the Verb Condition pointed more toward the action video ($M = .74$; $SEM = .05$) than those in the Noun Condition ($M = .30$; $SEM = .05$).

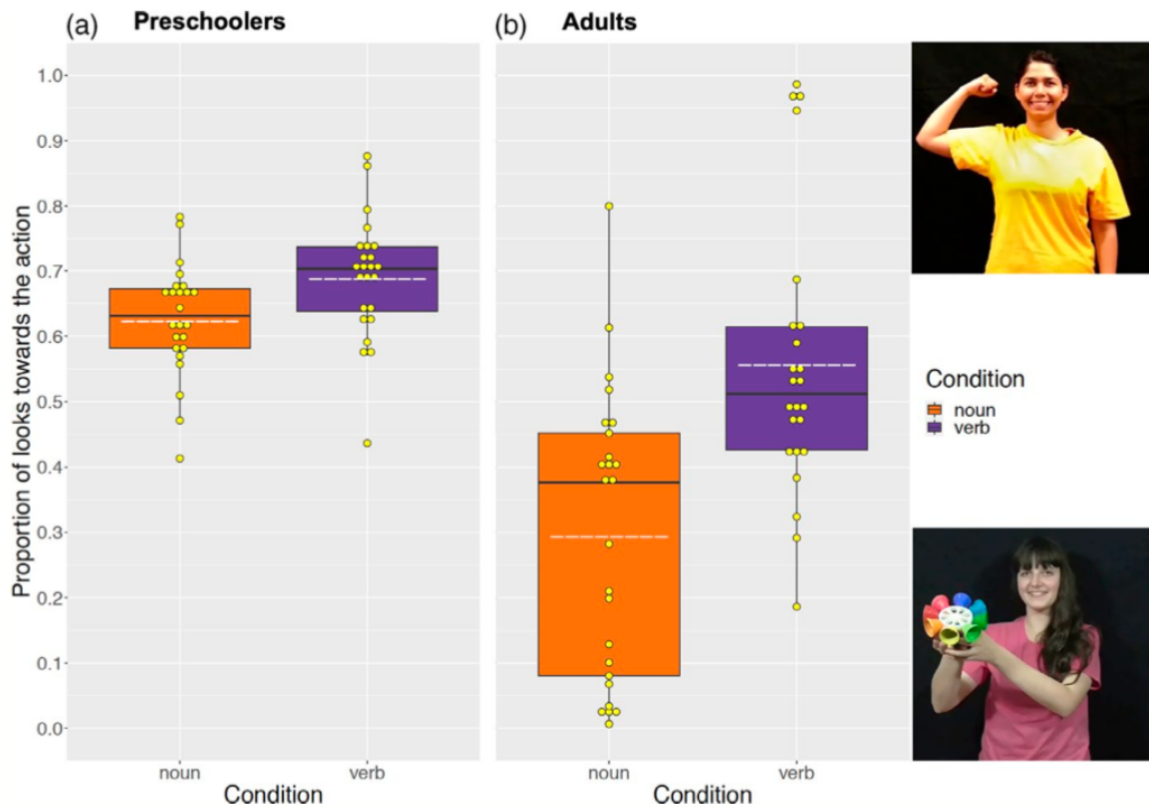


Figure 4: Preschoolers' (a) and adults' (b) mean overall proportion of looks toward the action video in the test phase for the two conditions (Verb Condition in purple and Noun Condition in orange). Yellow dots represent individual participants. The lower and upper hinges correspond to the first and third quartiles, the dotted white lines represent the means, and the black lines within the squares represent the median. The top whiskers denote the maximum value, and the bottom whiskers the minimum value. Note that outliers (one child in the Verb condition and another one in the Noun condition; four adults in the Verb condition) were not removed from any of the reported analyses. For each group, participants from the Verb condition are looking at the action video significantly more than those from the Noun condition

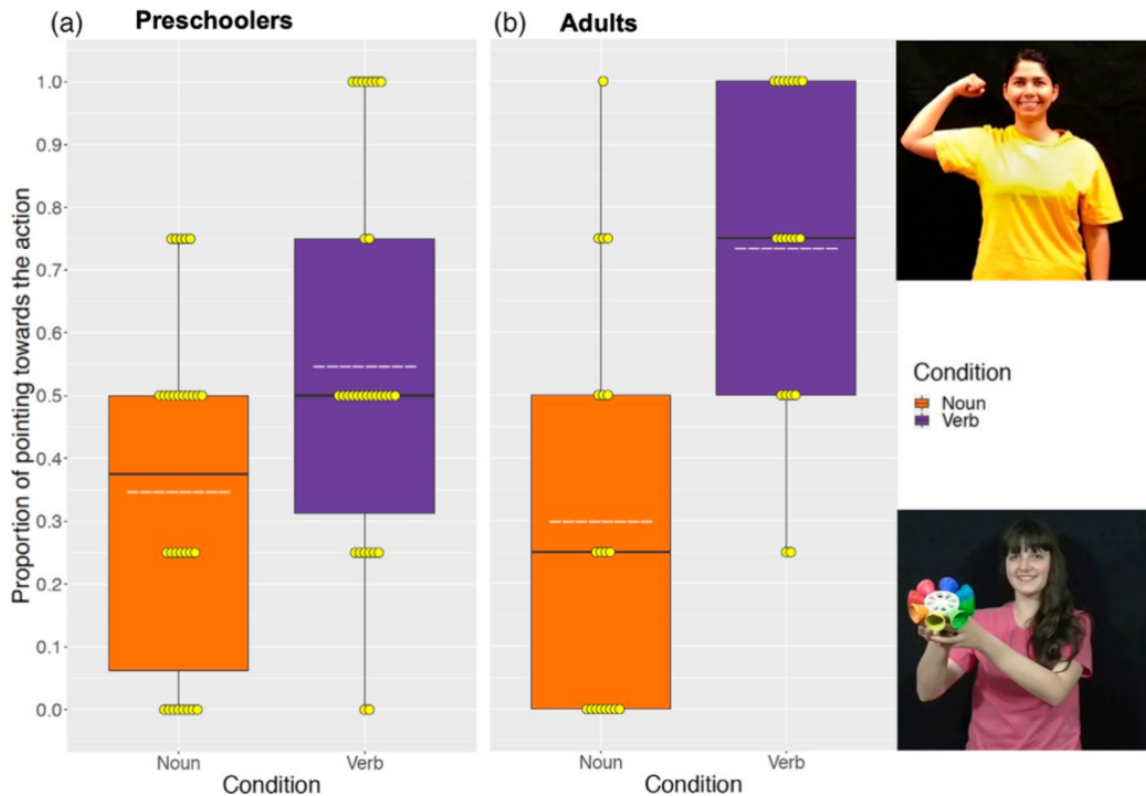


Figure 5: Boxplots of preschoolers' (a) and adults' (b) proportion of pointing towards the action. Each yellow dot represents one participant (with each score based on up to four novel target words). White dashed lines represent the means of the distributions

In summary, both looking behavior and pointing responses were significantly influenced by participants' assigned condition, with children in the Verb condition looking and pointing more towards the action video than children in the Noun condition. Adults' results were similar to those obtained with preschoolers. Overall, these results indicate that children in the Noun and Verb condition learned something different about the novel function word "ko". One possible interpretation is that children in the Verb condition learned that "ko" appears in verb contexts, and therefore interpreted a novel word following "ko" as more likely to refer to a novel action, while children in the Noun condition learned that "ko" appears in noun contexts, and thus at test they interpreted a novel word following "ko" as more likely to refer to a novel object. However, our results are also compatible with an alternative interpretation in which only one group of children learned something about the novel word "ko" (either the Noun

or the Verb group), and the other group was just baffled and did not learn anything. We cannot test this by comparing children's looking behavior to 'chance', since there is no guarantee that chance looking is at 50% -- indeed, our results show a bias with more looks towards the action video, in which there is distracting movement, compared to the object video¹. At any rate, the demonstration that the two groups of children learned something different about the novel function word "ko" shows that children of this age have the ability to track the use of a novel function word with familiar content words during a brief training phase and to use their knowledge later on to narrow down the probable meaning of co-occurring novel words, although it remains unclear whether children learned which semantic category was co-occurring with "ko" in both the Noun and the Verb conditions (or in only one of them).

Post-experiment questions to adult participants

After completing the experiment, a majority of the adult participants (46 out of 52; the total includes the three participants who were excluded based on their poor eye-tracking data) were able to explicitly describe the regularity (e.g. when they were in the Verb condition, they explained that "ko" was used as a pronoun or referred to the stuffed animal, and when they were in the Noun condition, that it was acting as a determiner or preceded things and animals), but they did not systematically report that their choices during the test phase (i.e. pointing/looking more towards the novel actions or towards the novel objects) was driven by this explicit knowledge. Many of them reported that they still had a doubt on how to use "ko" as a cue to novel word meaning. For instance, some participants in the Noun condition reported being aware that nouns can also denote events (e.g. a dance), not just objects.

¹ One reviewer suggested testing a control group of children who would not be exposed to a training video (and who would encounter 'ko' for the first time at test) could be an adequate measure of chance behavior. Such a control group is quite unpractical to run for two reasons: First, the difference in looking time between the two experimental groups is already rather small: if the control group fell in-between these two experimental groups, there is a good chance that its looking time would not be different from either experimental condition, making interpretation of the results difficult.

Discussion

Our study provides the first evidence that young children can rapidly track the distribution of a novel function word and learn what type of semantic properties it can predict. After watching a 5-min training video in which a novel function word “ko” either replaced pronouns (Verb condition) or articles (Noun condition), 3-4-year-olds were able to infer the probable meaning of novel content words that co-occurred with this newly-learned function word “ko”(e.g. *Oh look, ko dase!*), interpreting them as either referring to novel actions or novel objects. Children’s behavior during the test trials was most likely conditioned by their experience with “ko” during the training video. Children in the Verb condition heard “ko” preceding verbs that were referring to actions that were being performed by the agents, whereas children in the Noun condition heard “ko” preceding nouns that referred to the agents themselves (i.e. the stuffed animals) or to inanimate objects (e.g. the car). Since the training videos presented “ko” embedded in full sentences, and participants’ focus was on the storyline, children had to use their distributional analysis skills in order to keep track of the kind of words which co-occurred with the new function word. Hence, our results demonstrate that children at this age can proficiently use a newly-learned syntactic context as a “zoom lens” to guide their real-time interpretations of the meaning of novel words, mapping novel nouns to objects and novel verbs to actions.

The present findings have important implications for our understanding of the synergies between semantics and syntax during language development. While numerous studies over the past decades have focused on the use of syntactic contexts to guide young children’s acquisition of novel content word meanings (i.e., “they dax”, *dax* is a verb and likely refers to an action), a process called syntactic bootstrapping, the current study focuses on the process through which children might come to learn about the properties of specific syntactic contexts in their native language: it provides the first evidence that young children can exploit known content

words to learn the properties of a novel function word (i.e., “ko eats”, “ko plays”, “ko is reading” => “ko” is followed by action-denoting words). These results bring empirical evidence to the hypothesis that children expect words from similar conceptual categories to occur in similar syntactic contexts (e.g., Gleitman, 1990; Pinker, 1984), while also supporting the modeling-based learning mechanism of the *semantic seed* (e.g., Brusini, Amsili, Chemla, & Christophe, 2014; Christophe et al., 2016; Gutman et al., 2014). Our study shows that around 3-to-4 years of age, children rapidly and efficiently undertake a distributional analysis to build predictions about the type of concept that can co-occur with a newly-learned function word.

We would expect even younger children to rely on this learning strategy as well, since it would be most useful during the second year of life, when infants already have enough knowledge of the lexicon (e.g. Bergelson & Swingley, 2012), and the necessary computational skills (e.g. Gomez & Gerken, 1999). Although it would be fascinating to test it directly, our experimental task is too difficult for toddlers: even when they are tested with real determiners and pronouns (rather than newly-learnt ones) 20-month-olds have failed the online task of allocating their visual attention towards the appropriate meaning (action vs object) for novel content words (de Carvalho, Babineau, Trueswell, Waxman, and Christophe, 2019). Nevertheless, previous work has demonstrated infants’ ability to use frequent function words to build syntactic and semantic expectations shortly after their first birthday. For instance, at only 14 months of age, infants expect novel words to follow a determiner (and not a pronoun) if they were first encountered following other determiners (Babineau, Shi & Christophe, under review; Shi & Melançon, 2010; Höhle, Weissenborn, Kiefer, Schulz, & Schmitz; 2004). At 18 months of age, after an infant-controlled habituation phase, infants can map a novel word to either the action performed by an agent or the agent itself (the name of the animal) depending on the syntactic context in which this novel word appeared (e.g. “It is a bamoule” vs “She is bamouling”; de Carvalho, He, Lidz, & Christophe, 2019; He & Lidz, 2017). These latter

findings demonstrate that infants have acquired the two fundamental prerequisites for learning the semantic predictiveness of function words. In the present experiment, we simply took advantage of the 3-to-4-year olds' fast inferential processing skills in order to obtain direct evidence of the learning mechanism that bootstraps the syntactic bootstrapper.

It is worth mentioning that children's ease to integrate a novel function word in our study might have been facilitated by the chosen semantic categories (i.e. objects vs actions) co-occurring with it, which are already marked in their native language's morphosyntax. That is, French-speaking preschoolers already master a language in which verbs are preceded by personal pronouns and nouns by articles, so learning the predictiveness of the novel function word "ko" might have been facilitated by its familiar position and use during the training video. Moreover, it remains unclear whether children interpreted the novel function word "ko" as a synonym of the real function words in their language or as a new member for an already existing group of function words (i.e., a new kind of personal pronoun or article). In the control group with adults, a majority of participants reported that they were aware of such equivalents, although their generalization to novel words during the test phase was not at ceiling. Interestingly, a minority of adults reported to have no clue on the use or function of the novel function word. Therefore, implicitly learning the predictiveness of "ko" was a challenging task, even if this novel function word mimicked an already existing group of function words in the participants' native language. As a final point, we acknowledge that the mapping between nouns and objects, and between verbs and actions, is far from perfect. Obviously, many nouns also refer to actions (e.g. the dance, the construction), not only to concrete objects. This blurred line between the chosen distinctions might have added some noise, perhaps especially for adult participants who know many more action words belonging to the noun category.

In light of these limitations, future work should explore whether children can readily learn to predict other types of semantic categories that are not marked in their native language.

We suspect that learning a new type of syntactic marker would require a lengthier exposure (more than a 5-min video). It is still unclear whether there are limits on the types of lexical semantics that can be inferred through this *semantic seed* learning mechanism. Inspection of cross-linguistic universals reveals that across the world's languages, certain conceptual categories are often marked in morphosyntax (e.g. object/action, animate/inanimate distinctions), while others are typically not (e.g. electric appliance). As Strickland (2017) proposed, the cause of the cross-linguistic regularities might be that since morphosyntax helps infants to learn word meanings, only morphosyntactic regularities that mark conceptual distinctions that are noticeable by infants (e.g. *core knowledge*) would be selected by human languages. Overall, children's ability to pay attention to function words and their distribution, which starts in infancy and develops during early childhood, represents a powerful aid to learn a large lexicon without explicit teaching, which is one of the hallmarks of human languages.

Author Contributions

All authors made a significant contribution to the study. A. de Carvalho, J. Trueswell, and A. Christophe developed the initial study concept and contributed with M. Babineau to the study design. A. de Carvalho programmed the experiment and prepared the stimuli used during the test phase. Stimuli preparation for the training videos, testing, data collection and data coding were performed by M. Babineau. Statistical analyses were computed by M. Babineau. M. Babineau wrote the manuscript, and A. de Carvalho, J. Trueswell, and A. Christophe provided critical revisions. All authors approved the final version of the manuscript for submission.

Acknowledgements

This research was supported by a postdoctoral fellowship from the Fyssen Foundation to Mireille Babineau and to Alex de Carvalho and by a postdoctoral fellowship from the H2020 European Research Council [Marie Skłodowska-Curie grant No 799380) to Mireille Babineau. It was also supported by grants from the National Institute of Child Health and Human Development (Grant R01HD37507 to John Trueswell), Fondation de France, and the French ANR (ANR-13-APPR-0012 LangLearn, ANR-17-CE28-0007-01 LangAge, and ANR-17-EURE-0017 FrontCog).

The authors thank all the bright “actors” who participated in the videos, as well as Anne-Caroline Fiévet for her help in recruiting the children and organizing the partnerships with preschools, and Lara Narbona Sabaté who assisted us in the recording of the training videos. We also thank all the parents, and preschoolers who participated in this study and the directors and teachers of the preschools that so warmly welcomed us and allowed us to conduct our research.

Declaration of Conflicting Interests

The authors declare that they had no conflicts of interest.

References

- Arunachalam, S. (2016). Preschoolers' Acquisition of Novel Verbs in the Double Object Dative. *Cognitive Science*, 1–24. <https://doi.org/10.1111/cogs.12368>
- Bergelson, E., & Aslin, R. N. (2017). Nature and origins of the lexicon in 6-mo-olds. *Proceedings of the National Academy of Sciences*, 114, 12916-12921. <https://doi.org/10.1073/pnas.1712966114>
- Bergelson, E., & Swingle, D. (2012). At 6 – 9 months, human infants know the meanings of many common nouns. *Proceedings of the National Academy of Sciences*, 109, 3253-3258. <https://doi.org/10.1073/pnas.1113380109>
- Bernal, S., Lidz, J., Millotte, S., & Christophe, A. (2007). Syntax Constrains the Acquisition of Verb Meaning. *Language Learning and Development*, 3(33), 325–341. <https://doi.org/10.1080/15475440701542609>
- Brusini, P., Amsili, P., Chemla, E., & Christophe, A. (2014). *Simulation de l' apprentissage des contextes nominaux / verbaux par n-grammes*, In Proceedings of TALN 2014 (Volume 2: Short Papers). Marseille, France, July 2014, pp. 505-510. Association pour le Traitement Automatique des Langues.
- Chemla, E., Mintz, T. H., Bernal, S., & Christophe, A. (2009). Categorizing words using “frequent frames”: What cross-linguistic analyses reveal about distributional acquisition strategies. *Developmental Science*, 12(3), 396–406. <https://doi.org/10.1111/j.1467-7687.2009.00825.x>
- Christophe, A., Dautriche, I., de Carvalho, A., & Brusini, P. (2016). Bootstrapping the syntactic bootstrapper. In: J. A. Scott & D. Waughtal (Eds.), *Proceedings of the 40th Annual Boston University Conference on Language Development* (pp. 75–88). Boston, United States.
- Christodoulopoulos, C., Roth, D., & Fisher, C. (2016). An incremental model of syntactic

- bootstrapping. In *Proceedings of the 7th Workshop on Cognitive Aspects of Computational Language Learning* (pp. 38–43). Berlin, Germany.
- Connor, M., Gertner, Y., Fisher, C., & Roth, D. (2010). Starting from Scratch in Semantic Role Labeling. In *Acl 2010 - 48th Annual Meeting of the Association For Computational Linguistics, Proceedings of the Conference* (pp. 989-998).
- Connor, M., Fisher, C., & Roth, D. (2012). Starting from scratch in semantic role labeling: Early indirect supervision. In A. Alishahi, T. Poibeau, and A. Korhonen, editors, *Cognitive Aspects of Computational Language Acquisition*. Springer.
- de Carvalho, A., He, A. X., Lidz, J., & Christophe, A. (2019a). Prosody and Function Words Cue the Acquisition of Word Meanings in 18-Month-Old Infants. *Psychological Science*, 30(3), 319–332. <http://doi.org/10.1177/0956797618814131>
- de Carvalho, A., Babineau, M., Trueswell, J. C., Waxman, S. R., & Christophe, A. (2019b). Studying the Real-Time Interpretation of Novel Noun and Verb Meanings in Young Children. *Frontiers in Psychology*, 10, 1–16. <http://doi.org/10.3389/fpsyg.2019.00274>
- de Carvalho, A., Dautriche, I., & Christophe, A. (2017). Phrasal prosody constrains online syntactic analysis in two-year-old children. *Cognition*, 163, 67–79. <https://doi.org/10.1016/j.cognition.2017.02.018>
- Fisher, C., Gertner, Y., Scott, R. M., & Yuan, S. (2010). Syntactic bootstrapping. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1, 143–149. <https://doi.org/10.1002/wcs.17>
- Fisher, C., & Gleitman, L. R. (2002). Language acquisition. In H. Pashler & R. Gallistel (Eds.), *Steven's handbook of experimental psychology* (Volume 3:, pp. 445–496). New York: Wiley.
- Gerken, L. A., Wilson, R., & Lewis, W. (2005). Infants can use distributional cues to form syntactic categories. *Journal of Child Language*, 32(2), 249–268.

<https://doi.org/10.1017/S0305000904006786>

- Gillette, J., Gleitman, H., Gleitman, L., & Lederer, A. (1999). Human simulations of vocabulary learning. *Cognition*, *73*, 135–176. [https://doi.org/10.1016/S0010-0277\(99\)00036-0](https://doi.org/10.1016/S0010-0277(99)00036-0)
- Gleitman, L. (1990a). The Structural Sources of Verb Meanings. *Language Acquisition*, *1*(1), 3–55.
- Gleitman, L. (1990b). The structural sources of verbs' meaning. *Language Acquisition*, *1*(1), 3–55.
- Gleitman, L., Cassidy, K., Nappa, R., Papafragou, A., & Trueswell, J. (2005). Hard words. *Language Learning and Development*, *1*(1), 23–64. <https://doi.org/10.1038/scientificamerican0491-138>
- Gómez, R. L., & Gerken, L. (1999). Artificial grammar learning by one-year-olds leads to specific and abstract knowledge. *Cognition*, *70*, 109–135. [https://doi.org/10.1016/S0010-0277\(99\)00003-7](https://doi.org/10.1016/S0010-0277(99)00003-7)
- Gutman, A., Dautriche, M. I., Crabbé, D. B., & Christophe, D. A. (2014). Bootstrapping the Syntactic Bootstrapper: Probabilistic Labeling of Prosodic Phrases. *Language Acquisition*, *22*:3, 285-309. DOI: [10.1080/10489223.2014.971956](https://doi.org/10.1080/10489223.2014.971956)
- Hahn, N., Snedeker, J., & Rabagliati, H. (2015). Rapid Linguistic Ambiguity Resolution in Young Children with Autism Spectrum Disorder: Eye Tracking Evidence for the Limits of Weak Central Coherence. *Autism Research*, *8*(6), 717–726. <https://doi.org/10.1002/aur.1487>
- He, A. X., & Lidz, J. (2017). Verb Learning in 14- and 18-Month-Old English-Learning Infants. *Language Learning and Development*, *13*(3), 335–356. <https://doi.org/10.1080/15475441.2017.1285238>
- Höhle, B., Weissenborn, J., Kiefer, D., Schulz, A., & Schmitz, M. (2004). Functional

elements in infants' speech processing: The role of determiners in the syntactic categorization of lexical elements. *Infancy*, 5(3), 341–353.

https://doi.org/10.1207/s15327078in0503_5

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, 164(1), 177–190.

<https://doi.org/10.1016/j.jneumeth.2007.03.024>

Mintz, T. H. (2003). Frequent frames as a cue for grammatical categories in child directed speech. *Cognition*, 90(1), 91–117. [https://doi.org/10.1016/S0010-0277\(03\)00140-9](https://doi.org/10.1016/S0010-0277(03)00140-9)

Shi, R., & Melançon, A. (2010). Syntactic Categorization in French-Learning Infants.

Infancy, 15(5), 517–533. <https://doi.org/10.1111/j.1532-7078.2009.00022.x>

Strickland, B. (2017). Language Reflects “Core” Cognition: A New Theory About the Origin of Cross-Linguistic Regularities. *Cognitive Science*, 41(1), 70–101.

<https://doi.org/10.1111/cogs.12332>

Tincoff, R., & Jusczyk, P. W. (1999). Some beginnings of word comprehension in 6-month-olds. *Psychological Science*, 10(2), 172–175. <https://doi.org/10.1111/1467-9280.00127>

Von Holzen, K., & Mani, N. (2012). Language nonselective lexical access in bilingual toddlers. *Journal of Experimental Child Psychology*, 113, 569–586.

<https://doi.org/https://doi.org/10.1016/j.jecp.2012.08.001>

Weisleder, A., & Waxman, S. R. (2010). What's in the input? Frequent frames in child-directed speech offer distributional cues to grammatical categories in Spanish and English. *Journal of Child Language*, 37(5), 1089–1108.

<https://doi.org/10.1017/S0305000909990067>