Gathering Clues to Meaning: Children Recruit Past Linguistic Context to Infer Word Meanings Across Exposures in a Discourse

Idella Smolyar (ismolyar@usc.edu) & Alexander LaTourrette
Department of Psychology, University of Southern California

Abstract

Children's word learning often occurs under referential ambiguity, but linguistic context can drastically reduce this ambiguity and facilitate accurate word learning. Here, we examined whether linguistic context also facilitates word learning *across* exposures, testing 3-year-olds' ability to use a word's prior linguistic context to constrain subsequent mappings. To test this, we conducted three studies using semantically informative verbs, comparing children's use of linguistic context to infer meanings within a single exposure (Study 1 & 3) or across exposures (Study 2 & 3). We also varied the noun's salience in the discourse context between exposures. Results indicated that children successfully used familiar verbs to infer novel noun meanings both when the verb and referent co-occurred within an exposure and when they were separated across exposures. However, children were less successful when linguistic and referential contexts were separated. Thus, 3-year-olds use prior linguistic contexts to infer word meanings as a conversation progresses.

Keywords: word learning; development; linguistic context; discourse cues

Introduction

Infants' early word learning is relatively slow and difficult, with infants typically recognizing only a few words at 6 months of age (Bergelson & Swingley, 2012) and 50 to 100 by their first birthday (Frank et al., 2017). This slow growth is likely due in part to the challenge of referential ambiguity: when children encounter a new word, there are often multiple candidate referents for it (Medina et al., 2011). Indeed, children interacting with caregivers have an average of around eight object types in their view (Clerkin et al., 2017). Similarly, eye-tracking studies suggest that while children may only fixate on one or two objects after hearing a word, the objects they fixate on are no more likely to be the target object than irrelevant objects (Yu et al., 2021). Thus, early word-referent mappings are likely quite error-prone.

As children's vocabularies grow, however, they can rely on a new word's linguistic context to reduce this referential ambiguity and facilitate accurate word learning (Goodman et al., 1998; Yuan & Fischer, 2009). For example, hearing "Let's *eat* the *snickerdoodles*" (rather than "We've *got* the *snickerdoodles*") might help a learner narrow down the set of plausible referents for "snickerdoodle" (i.e., to food items). Past work has shown children can leverage a range of linguistic cues to infer word meanings. This includes using known adjectives to infer the meanings of both nouns ("The dax is happy": Syrett et al., 2019) and adjectives ("The chromium one, not the blue one": Carey & Bartlett, 1978). Similarly, children use known nouns and word order to infer verb meanings ("The duck is gorping the bunny": Gertner et al, 2006) and use transitive frames to infer verb meanings ("She blicks the baby": Arunachalam & Waxman, 2010; Naigles, 1990; Yuan & Fisher, 2009). A robust series of findings also shows that 2- to 3-year-old children use familiar verbs to infer the meanings of novel nouns introduced as the verb's subject (Ferguson et al., 2014, 2018) or object (Borovsky et al., 2016, Goodman et al., 1998; LaTourrette et al., 2023). In each of the above cases, children use linguistic context to select the target referent from an assortment of novel objects on an immediate test – e.g., hearing "Mommy feeds the dax" and picking the sole animal as the feedable referent. These meanings are even successfully retained after 4- to 24-hour delays (Goodman et al., 1998; He et al., 2020).

However, most experimental paradigms are carefully constructed to provide children with the linguistic and referential context either simultaneously (e.g., Carey & Bartlett, 1978; Goodman et al., 1998; LaTourrette et al., 2023), or immediately following each other with no delay (e.g., Ferguson et al., 2014, Scott & Fisher, 2009). Similarly, children utilize linguistic context from an immediately preceding sentence in a discourse to infer the meaning of a new word (Sullivan & Barner 2016, 2019). In everyday life, on the other hand, children will often hear words without the target referent being immediately evident (Yu et al., 2021). Indeed, we often talk about absent objects. For example, a child may hear that they will "eat a snickerdoodle for dessert" or "feed a llama at the zoo later", with the noun's referent not appearing until after at least a brief delay. Words are often uttered in the absence of their referents (Cartmill et al., 2016), including both nouns (Medina et al., 2011) and verbs (Tovar-Perez et al., 2025). Moreover, even if the target is present, children often fail to identify it (Yu et al., 2021). It is a crucial question, then, whether linguistic context still supports word learning *across* exposures—when the linguistic and referential context do not co-occur.

Promising evidence for the use of linguistic context across exposures comes from children's learning of transitive verbs.

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

Hearing a verb used in a transitive frame presented in the absence of any referential events (e.g., "Yes, she blicked the cat") led 2-year-olds to preferentially map that verb to a causative event over a non-causative event, and they were similarly successful even after a 24–48-hour delay (Yuan & Fisher, 2009; Arunachalam & Waxman, 2010). That said, this kind of cross-exposure inference has so far only been shown with causal verbs. For instance, while transitive frames lead children to map verbs to contact events on an immediate test, this is not true after a delay (Arunachalam & Dennis, 2019). Thus, questions remain about whether and how well cross-exposure effects of linguistic context generalize to different types of words.

Current Work

In the present work, we directly addressed these questions, testing children's ability to use linguistic context across exposures when inferring noun meanings. In Study 1, we tested the influence of linguistic context *within* a single exposure, extending prior work. Then in Study 2, we investigated the influence of linguistic context *across* exposures embedded in a story discourse. For Study 2, children heard the same linguistic context as in Study 1, but in the absence of a referent. We then tested if they used this context to infer the noun's meaning from among multiple referents on a later exposure. Lastly, for Study 3, we ran a conceptual replication of Studies 1 and 2 using a broader set of counter-balanced stimuli, comparing the use of linguistic context to infer noun meanings within- or across-exposures.

Study 1

First, we examined children's use of linguistic context within a single exposure, presenting the informative linguistic context (e.g., "I want to *read* the dax") simultaneously with the possible referents (e.g., an image of a book and of a bike). Three-year-olds learned 5 novel nouns, each paired with a verb (e.g., "read," "feed," "eat") found to support semantic inferences in prior work (Borovsky et al., 2016; Goodman et al., 1998; LaTourrette et al., 2023).

Methods

Participants We ran all studies on the Children Helping Science platform. For Study 1, we recruited 26 3-year-olds (15 females, 11 males, M_{age} = 3.55, SD = .31), with another 4 children excluded for below-chance performance on familiar word trials (n=2) or not completing the study (n=2). The study took 10 minutes, and all parents gave informed consent and were compensated with $5 gift cards. We expected an effect size of d = 1, similar to that reported for 2-year-olds using linguistic context to infer verb meanings (Arunachalam et al., 2016; Yuan & Fisher, 2009). At an alpha level of .05, a sample size of 16 yielded 95% power, but we used a 50% larger sample to allow for noise in online testing.

Materials On each trial, children viewed two objects (e.g., a book and a bicycle) and heard two sentences using an informative verb compatible with only one of the objects (e.g., "I like to read the dax! Yes, I read the dax!"). Children were then prompted to point to the novel noun's referent (e.g., "Point to the dax!"). Children were exposed to 5 novel words and 13 familiar words, used as filler trials. All audio materials were recorded by a female native English speaker in child-directed speech. The informative verbs used in novel trials included: "climb", "eat", "feed", "read", and "wear". Visual materials consisted of 40 pictures of objects from familiar categories. We used familiar objects as the referents of novel words to ensure that children understood the affordances of objects (e.g., after hearing "I want to read the dax," they could readily identify the readable object).

Procedure During the task, children were exposed to 13 "filler" trials using familiar nouns, 5 trials using novel nouns, and another 8 trials for a different task norming other stimuli. Familiar word trials always intervened between novel word trials. On each trial, children were presented with 2 objects and prompted to select one of them. On filler trials, audio prompts featured an uninformative verb and familiar noun (i.e., "I see a cow! Point to the cow!"). Novel noun trials featured an informative verb paired with a novel noun (i.e., "I like to read the dax! Point to the dax"). In each novel noun trial, the verb was compatible with only one of the two visible referents. Thus, children always chose between a verb-compatible and verb-incompatible referent.

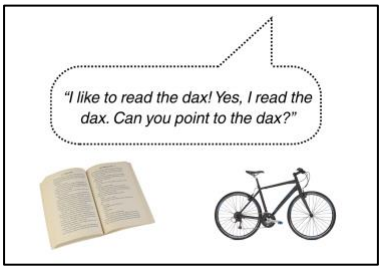

Figure 1. Study 1: Sample Novel Noun Trial. Children viewed two images, heard a sentence using an informative verb compatible with only one object, and were then asked to select the novel word's referent.

Results

First, we examined children's performance on the "filler" trials featuring familiar nouns. As expected, children performed at near-ceiling levels (M = .97, SD = .03), indicating they were highly engaged with the task.

Next, to test whether children used the informative verbs to infer noun meanings within an exposure, we constructed an intercept-only binomial mixed-effects model with a random effect of participant and item, predicting children's selection of the verb-compatible referent. This model indicated that children preferentially selected the verb-compatible referent (M = .82, SD = .21, d = 1.52), doing so at an above-chance rate (β = 1.74, SE = .36, z = 4.89, p < .001). Notably, children scored above chance on each of the verbs, indicating robust learning. Thus, three-year-olds successfully used a noun's linguistic context within an exposure to select the target, verb-compatible referent. Finally, we tested for an age effect, adding a fixed effect of age to the model. This yielded a non-significant effect of age (β = -.013, SE = .01, z = -1.05, p = .30).

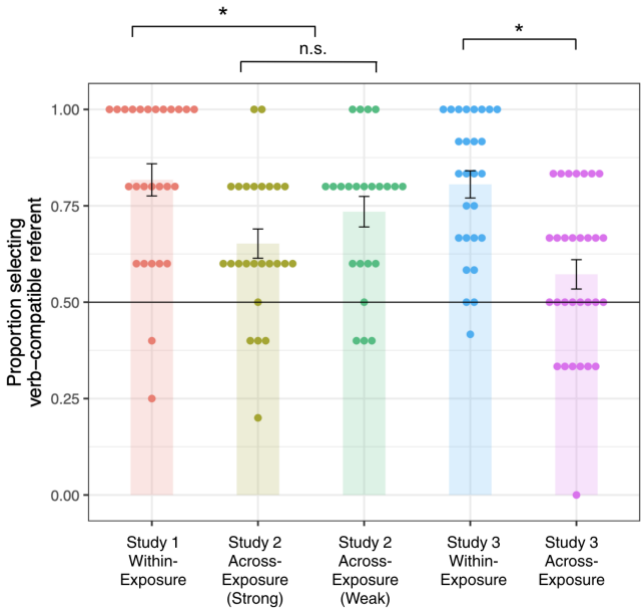

Figure 2. Results for Studies 1-3: Children selected the verb-compatible referent more when the verb co-occurred with the referent (Study 1 & 3, Within-Exposure) than when separated (Study 2 & 3, Across-Exposure).

Study 2

In Study 1, we found that children successfully used informative verbs to infer novel nouns' referents. But as noted above, a noun's referent is often not present in the child's view when the noun is uttered (Yu et al., 2021). Thus, in Study 2, we examined whether children could also use the informative linguistic contexts from Study 1 to infer noun meanings *across* exposures. Specifically, we presented children with a discourse in which the informative linguistic context (e.g., "I want to read a dax") was encountered with no referents present. This Learning exposure was then followed by discussion of other objects in the conversation, before the speaker returned to the word on the Test trial but without pairing the noun with its informative verb again ("There's the dax!"). As in Study 1, children then choose between two objects as the referent of the noun – with only one of the referents being compatible with the previously heard informative verb. The key question was whether children would leverage the previous exposure's linguistic context to successfully infer the noun's referent later in the conversation.

Past work has shown that children can use some discourse structure cues to infer word meanings. However, like other work on children's use of linguistic context (e.g., Ferguson et al., 2018), this work only looked at inferences from the immediately preceding sentence, not inferences across exposures. For instance, work from Sullivan et al. (2019) showed that if children hear a sentence like "I'm really thirsty; I wish I had a dax", children tend to map "dax" to a drink – using the current discourse context to infer a new word's meaning. Massing exposures to a word within a discourse can also help word learning (Schwab & Lew-Williams, 2016). However, the current paradigm asks children to go further, drawing on information from a previous exposure to infer the word's meaning across a short delay instead of just in-the-moment.

Following past work on the role of discourse context, we also manipulated the structure of our discourse, creating Strong and Weak Discourse Contexts. In both cases, children met a speaker who introduced a novel noun paired with an informative verb but in the absence of a referent (e.g., "I want to read a dax. Let's go find a dax!"). In the Strong Discourse Context condition, the speaker then continued to mention the novel noun, making it the focus of the conversation. This emulates the bursty structure of kids' input, where nouns often get mentioned many times in a particular burst (Slone et al., 2023). We compared that structure to a Weak Discourse Context condition, where children did not hear the novel noun between the initial, linguistically informative Learning exposure and the Test trial.

Methods

Participants Children were recruited and completed the study just as in Study 1. We recruited 48 3-year-olds (26 females, 22 males, M = 3.44 years, SD = .31), with an additional 4 participants excluded for performing below chance on familiar "filler" trials (n = 1) or didn't complete the full study (n = 3). Based on our effect size from Study 1 (d = 1.52), we recruited 24 kids per condition, providing 99% power (α = .05).

Materials We included 36 total images for the targets and distractors, all identical to Study 1. However, additional audio was included: at the start of the study, a narrator introduced 3 characters (2 female, 1 male). Each character introduced two novel words, in separate stories. We used the

same 5 novel informative verbs from Study 1 but added an additional verb ("pet") for counterbalancing across narrators. Past work suggests that 3-year-olds successfully use "pet" to map its object to (furry) animals (Borovsky et al., 2016). Note that all patterns of significance are identical, regardless of whether the "pet" item is included in the analyses, so we report the full analyses here.

Design Children were placed in one of two conditions: the Strong Discourse Context (speakers reminded children of the novel noun after each filler trial between Learning and Test) or Weak Discourse Context (speakers did not mention the noun between Learning and Test exposures). Note that both conditions presented the words embedded in a short, ongoing, goal-oriented discourse with a single character.

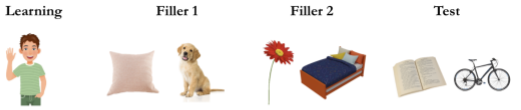

Condition	Learning	Filler 1	Filler 2	Test
Strong Discourse Context	I want to **read** a dax. Let's go **read** a dax!	**We're on our way to find a dax.** Look! There's a dog. Point to the dog.	**We're almost to the dax.** Wow! There's a flower. Point to the flower.	There's the dax! Point to the dax!
Weak Discourse Context	I want to **read** a dax. Let's go **read** a dax!	Look! There's a dog. Point to the dog.	Wow! There's a flower! Point to the flower.	There's the dax! Point to the dax!

Figure 3. Study 2 Design: At Learning, children heard a novel noun introduced with an informative verb. After two Filler trials, children were prompted to select the novel word's referent at Test, with only one of the two possible referents being compatible with the prior verb.

Procedure Similar to Study 1, children completed a series of trials featuring 12 familiar nouns and 6 novel nouns. However, for each novel noun, instead of encountering the possible referents and the informative verb simultaneously, these were now separated across exposures.

On the Learning exposure, children heard a character use the novel noun twice with an informative verb (e.g., "Let's go read a channer!"). On this exposure, there were no referents visible on the screen; all objects were introduced as goals that the character wanted to find, a natural context for mentioning nouns without their referents present *(see Fig. 3)*. Following each Learning exposure, children completed two filler trials. On each filler trial, the character prompted children to select the object corresponding to a familiar noun (e.g., "Find the cow!"). In the Strong Discourse Context condition, the speaker would mention the novel noun at the beginning of each filler trial, before prompting children to choose among the current objects (e.g., "We're almost to the channer. Oh look! There's a flower; point to the flower!"). In the Weak Discourse Context, the novel noun was not mentioned on filler trials.

After the filler trials, all children then arrived at the Test trial, where they saw two possible referents (identical to Study 1) and were asked to point to the target noun's referent. The time between Learning and Test trials varied depending on how long children took to select the familiar nouns' referents, but it generally represented only a 45-60 second delay. Thus, if 3-year-olds can use a word's prior linguistic context to infer its meaning, they should select the verb-compatible referent at Test, as they did in Study 1.

Results

First, we examined children's performance on the "filler" trials featuring familiar nouns. As expected, children performed well (M = .91, SD = .13), in line with our results from Study 1.

Across Exposures We first constructed a mixed-effects model for each condition, predicting whether children selected the verb-compatible referent. These included random effects of participant and item. A mixed-effects regression combining data from both conditions and including a sum-coded fixed effect of condition found reliably above-chance performance across conditions, (β = .87, SE = .20, z = 4.29, p < .001) with no difference between the Strong Discourse Context (M = .65, SD = .19, d = .79) and Weak Discourse Context (M = .73, SD = .21, d = 1.09), (β = .24, SE = .28, z = .83, p = .40) *(see Fig. 2)*. This effect was also consistent across items, with all verbs being at or above 50%. Finally, we tested for an age effect, adding a fixed effect of age to the model. This yielded a non-significant, negative effect of age (β = -.014, SE = .007, z = -1.89, p = .06). In sum, children in both the Strong and Weak condition significantly preferred verb-compatible referents, with no apparent benefit of the "bursty" mentions of the noun in the Strong Discourse Context.

Within vs. Across Exposures To test if the separation of linguistic and referential contexts affected learning, we then compared Study 1 and 2 *(Fig. 2)*. We constructed a binomial mixed-effects model predicting selection of verb-compatible referents, including random effects of participant and item and a sum-coded fixed effect of Study. Results indicated children chose the verb-compatible referent more often when the verb co-occurred with the referent (Study 1) than when separated (Study 2), (β = .81, SE = .33, z = 2.48, p = .013). Thus, while 3-year-olds performed above chance in all cases, they were significantly less successful in making semantic inferences across separate exposures than within an exposure.

Study 3

The results of Study 2 suggest that while children have more difficulty leveraging linguistic context across exposures, they can still do so. To strengthen this conclusion, and provide a rigorous comparison against chance, we ran an additional, better counterbalanced design in Study 3. In Studies 1 and 2,

the target and distractor (non-target) objects were not counterbalanced, so it is possible that the target objects could have unintentionally been more salient or attractive to children. To address this concern, we designed a follow-up fully counterbalancing the target and distractor objects, so that across children, each object served as both a target and a distractor item. For example, in Studies 1 and 2, the distractor on "read" trials was always a bike (whereas the bike was never the target on another trial). However, in Study 3, the distractor on "read" trials was an animal (e.g., a fish) that would be the target on another child's "feed" trial, where the book would now serve as a distractor *(see Fig. 4).* We also added an additional object per target noun, ensuring each child would see a different target/distractor pair for each of the six verbs. Thus, in some cases, the target object of "read" would be a list, while the distractor was a bunny.

Using this counterbalanced structure, we compared a Within-Exposure Condition, as in Study 1, against an Across-Exposure Condition, as in Study 2.

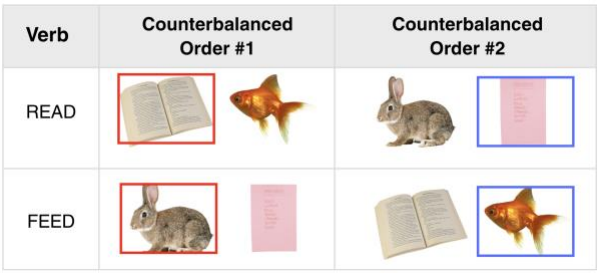

Figure 4. Study 3 Sample Test Trials, featuring 2 of the 6 verbs used in two counterbalanced orders. On each Test trial, children viewed two images, with the distractor object for one verb (i.e., the fish for "read" in Counterbalanced Order 1) serving as the target object for another verb in the other order (i.e., becoming the target of "feed" in Counterbalanced Order 2). Here, the target object is highlighted in the colored square *(the square was not seen by participants)*.

Methods

Participants For Study 3, a total of 57 children participated, with 27 in the within-exposure Study 1 replication (15 females, 12 males) and 30 (18 females, 12 males) in the across-exposure Study 2 replication. The studies looked at 3-year-olds once more: the Within-Exposure condition had a mean age of 3.43 years (n = 27, *SD* = .29) and the Across-Exposure condition of 3.49 years (n = 30, *SD* = .33). Based on our effect size from Study 2 (d = .95), we expected this would yield 99% power to detect such an effect (α = .05).

Materials All audio were identical to Study 1 and 2, along with number of trials per experiment. However, due to our new counterbalancing measures, we now included two target objects compatible with each verb. To facilitate counterbalancing, we also replaced the verb "pet" with "open," (animate referents were difficult to counterbalance as they were plausibly compatible with multiple other verbs, e.g., "eat"). This resulted in a total of 6 novel noun/verb pairs and 12 possible referents, each of which served as a target and non-target equally often across children.

Design Participants were assigned to one of two between-subjects conditions: a Within-exposure condition similar to Study 1 or an Across-exposures condition similar to Study 2's Strong Discourse Context. Since there was no significant difference between conditions in Study 2, we decided to provide children with the task that was in-principle easier.

Procedure The procedure was similar to Study 2, including 12 trials featuring familiar nouns interleaved with 6 trials featuring novel nouns paired with an informative verb. In the Within-Exposure condition, children followed the same procedure as Study 2 except that instead of separate Learning and Test trials, there was only a Test trial, which featured both the informative linguistic context and the referents, as in Study 1 (e.g., "I want to read a dax! Let's go read a dax. Point to the dax!"). The Across-Exposure condition was identical to Study 2's Strong Discourse Context condition.

Results

Within Exposures We first examined performance in our Within-Exposure condition, replicating Study 1. Children performed near ceiling on filler trials (M = .97, SD = .10).

We also observed above-chance performance on the novel noun trials in the Within-Exposure condition: children preferentially selected the verb-compatible referents (M = .81, SD = .11). This preference was reliably above chance, as indicated by a significant intercept term in a mixed-effects model predicting verb-compatible selections with random intercepts for participant and target (β = 1.94, SE = .38, z = 5.04, p < .001). We then compared performance in Study 1 and 3. A mixed-effects model with a sum-coded fixed effect of Study and random intercepts for participant and target showed no significant difference between studies (β = -0.06, SE = .53, z =-0.12, p = .91). Therefore, with a new, balanced set of stimuli, we replicated our Study 1 finding that 3-year-olds successfully use a noun's linguistic context to identify its referent within a single exposure.

Across Exposures For our Across-Exposure condition, we first examined results in the filler trials and found that children performed well (M = .99, SD = .03). To examine performance on Test trials, we constructed a binomial mixed-effects model predicting selection of verb-compatible referents, including random effects of participant and item. The results indicated that 3-year-olds still preferentially selected verb-compatible referents (M = .57, SD = .21, d = .27) and did so at above chance-rates (β = .34, SE = .15, z = 2.25, p = .024), confirming our earlier findings. Indeed, a

mixed-effects model comparing Study 2's Strong Discourse Context condition with the present results yielded a non-significant fixed effect of Study, ($\beta = .83$, $SE = .60$, $z = 1.38$, $p = .166$). However, we did observe greater variability across verbs than in Study 2, with only 4/6 verbs performing above chance in Study 3.

Within vs. Across Exposures Finally, we compared our Study 3 Within- and Across-Exposure conditions. To do so, we constructed a binomial mixed effects model predicting verb-compatible selections. Using a fixed effect of Condition and random effects for child and target word, we found that children were less likely to select the verb-compatible target image in the Across-Exposure condition compared to the Within-Exposure condition ($\beta = -1.37$, $SE = .32$, $z = -4.31$, $p < 0.001$) *(see Fig. 2).* Finally, we tested for an age effect, adding a fixed effect of Age to the model. This yielded a non-significant effect of Age ($\beta = -.002$, $SE = .01$, $z = -.18$, $p = .86$). Thus, 3-year-olds consistently used linguistic context to constrain referent mappings, but they were more successful when the linguistic context and referents co-occurred.

General Discussion

In three studies, we investigated the role of linguistic context in guiding word learning within and across exposures. Across all our studies, children consistently selected verb-compatible referents at above-chance rates—indicating that children successfully used familiar verbs to infer novel noun meanings both within a single exposure (Studies 1 & 3) and across exposures (Studies 2 & 3). The within-exposure results are in line with prior work (e.g., Borovsky et al., 2016; Goodman et al., 1998) and illustrate the robustness of such inference effects (effect sizes $d > 1$). As we saw little effect of age, these abilities appear to be robust by 3 years of age.

Moreover, this work provides the first evidence that children can leverage linguistic context to infer noun meanings *across* exposures as well. When multiple exposures are embedded within an ongoing discourse, children successfully use prior linguistic context to infer word meanings as the conversation progresses. This finding was evident in both Study 2 and the more stringently counterbalanced Study 3. This provides some of the first clear evidence outside transitive verbs (cf. Yuan & Fisher, 2009) that encountering a word in an informative linguistic context on one exposure will continue to guide children's semantic mappings on exposures after a delay.

That said, unlike for transitive verbs, children were significantly less successful when the informative linguistic and referential contexts were separated across exposures rather than co-occurring within an exposure. This weaker effect, especially in the more stringent structure of Study 3 ($d = .27$), suggests some limits on children's ability to retain or leverage linguistic context from prior exposures. This is especially true as our study had a relatively short delay of around 45 seconds between the exposures, with only familiar objects discussed in the interim—so children did not need to learn multiple words simultaneously. Despite this modest delay, children's performance significantly declined. It will be crucial, then, for future work to consider longer delays, like the 24–48-hour delays shown to be effective for transitive verbs in Yuan and Fisher (2009). Leveraging a word's linguistic and referential contexts across exposures separated by days or weeks would represent an even more powerful learning strategy, as illustrated by the success of language models in building semantic representations from a word's pattern of occurrences (e.g., Kwantes, 2005). However, the extent to which children rely on such memory-intensive, cross-exposure strategies is not yet clear.

Future work might also consider assessing less supportive conversational contexts: for instance, our informative linguistic contexts were always embedded in statements about the speaker's future goal ("I want to read a dax"). Encountering the same linguistic contexts in less goal-oriented discourses could yield weaker learning. Children's inferences were robust, however, even without reminders of the novel noun between Learning and Test exposures: performance in Study 2's Weak Discourse Context was well above chance and not significantly different from the Strong Discourse Context. Thus, cross-exposure inferences are robust to apparent changes in the topic of a conversation. However, our discourses were always conducted by a single speaker, who introduced and tested the same novel word. Future work should examine if children also draw generalizable semantic inferences across multiple speakers (Bulgarelli & Bergelson, 2023; Srinivasan et al., 2019).

Finally, it will be crucial to test children's cross-exposure inferences for other lexical categories, including verbs and adjectives. Future work might also consider a wider range of semantic cues to identify which cues children rely on most – for instance, directly comparing the use of transitive syntax to infer verb meanings with the current task's use of informative verbs to infer noun meanings. This would provide additional insight into the syntactic and semantic cues that are reliably retained and used by children to identify a word's correct meaning across delays.

In sum, the current results provide new evidence that children can successfully use linguistic context to infer noun meanings across exposures – enabling children to learn more from referentially ambiguous exposures as a conversation progresses. This integration of linguistic and referential information represents a potentially powerful word learning strategy. However, children's weaker performance across exposures (compared to within a single exposure) also raises intriguing new questions about the potential limits of this ability and the mechanisms underlying early word learning.

References

Arunachalam, S., & Waxman, S. R. (2010). Meaning from syntax: Evidence from 2-year-olds. *Cognition*, *114*(3), 442–446. https://doi.org/10.1016/j.cognition.2009.10.015

Arunachalam, S., & Dennis, S. (2019). Semantic detail in the developing verb lexicon: An extension of Naigles and Kako (1993). *Developmental Science*, *22*(1), e12697. https://doi.org/10.1111/desc.12697

Bergelson, E., & Swingley, D. (2012). At 6–9 months, human infants know the meanings of many common nouns. *Proceedings of the National Academy of Sciences*, *109*(9), 3253–3258. https://doi.org/10.1073/pnas.1113380109

Borovsky, A., Ellis, E. M., Evans, J. L., & Elman, J. L. (2016). Semantic structure in vocabulary knowledge interacts with lexical and sentence processing in infancy. *Child Development*, *87*(6), 1893–1908. https://doi.org/10.1111/cdev.12554

Bulgarelli, F., & Bergelson, E. (2023). Talker variability is not always the right noise: 14 month olds struggle to learn dissimilar word-object pairs under talker variability conditions. *Journal of Experimental Child Psychology*, *227*, 105575. https://doi.org/10.1016/j.jecp.2022.105575

Carey, S., & Bartlett, E. (1978). *Acquiring a single new word*.

Cartmill, E. A. (2016). Mind the gap: Assessing and addressing the word gap in early education. *Policy Insights from the Behavioral and Brain Sciences*, *3*(2), 185–193. https://doi.org/10.1177/2372732216657565

Clerkin, E. M., Hart, E., Rehg, J. M., Yu, C., & Smith, L. B. (2017). Real-world visual statistics and infants' first-learned object names. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *372*(1711), 20160055. https://doi.org/10.1098/rstb.2016.0055

Ferguson, B., Graf, E., & Waxman, S. R. (2014). Infants use known verbs to learn novel nouns: Evidence from 15- and 19-month-olds. *Cognition*, *131*(1), 139–146. https://doi.org/10.1016/j.cognition.2013.12.014

Ferguson, B., Graf, E., & Waxman, S. R. (2018). When *veps* cry: Two-year-olds efficiently learn novel words from linguistic contexts alone. *Language Learning and Development*, *14*(1), 1–12. https://doi.org/10.1080/15475441.2017.1311260

Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2017). Wordbank: An open repository for developmental vocabulary data. *Journal of Child Language*, *44*(3), 677–694. https://doi.org/10.1017/S0305000916000209

Gertner, Y., Fisher, C., & Eisengart, J. (2006). Learning words and rules: Abstract knowledge of word order in early sentence comprehension. *Psychological Science*, *17*(8), 684–691. https://doi.org/10.1111/j.1467-9280.2006.01767.x

Goodman, J. C., McDonough, L., & Brown, N. B. (1998). The role of semantic context and memory in the acquisition of novel nouns. *Child Development*, *69*(5), 1330. https://doi.org/10.2307/1132269

He, A. X., Huang, S., Waxman, S., & Arunachalam, S. (2020). Two-year-olds consolidate verb meanings during a nap. *Cognition*, *198*, 104205. https://doi.org/10.1016/j.cognition.2020.104205

LaTourrette, A., Chan, D. M., & Waxman, S. R. (2023). A principled link between object naming and representation is available to infants by seven months of age. *Scientific Reports*, *13*(1), 14328. https://doi.org/10.1038/s41598-023-41538-y

Medina, T. N., Snedeker, J., Trueswell, J. C., & Gleitman, L. R. (2011). How words can and cannot be learned by observation. *Proceedings of the National Academy of Sciences*, *108*(22), 9014–9019. https://doi.org/10.1073/pnas.1105040108

Naigles, L. (1990). Children use syntax to learn verb meanings. *Journal of Child Language*, *17*(2), 357–374. https://doi.org/10.1017/S0305000900013817

Schwab, J. F., & Lew-Williams, C. (2016). Repetition across successive sentences facilitates young children's word learning. *Developmental Psychology*, *52*(6), 879–886. https://doi.org/10.1037/dev0000125

Scott, R. M., & Fisher, C. (2009). Two-year-olds use distributional cues to interpret transitivity-alternating verbs. *Language and Cognitive Processes*, *24*(6), 777–803. https://doi.org/10.1080/01690960802573236

Slone, L. K., Abney, D. H., Smith, L. B., & Yu, C. (2023). The temporal structure of parent talk to toddlers about objects. *Cognition*, *230*, 105266. https://doi.org/10.1016/j.cognition.2022.105266

Srinivasan, M., Foushee, R., Bartnof, A., & Barner, D. (2019). Linguistic conventionality and the role of epistemic reasoning in children's mutual exclusivity inferences. *Cognition*, *189*, 193–208. https://doi.org/10.1016/j.cognition.2019.04.001

Sullivan, J., & Barner, D. (2016). Discourse bootstrapping: Preschoolers use linguistic discourse to learn new words. *Developmental Science*, *19*(1), 63–75. https://doi.org/10.1111/desc.12289

Sullivan, J., Boucher, J., Kiefer, R. J., Williams, K., & Barner, D. (2019). Discourse coherence as a cue to reference in word learning: Evidence for discourse bootstrapping. *Cognitive Science*, *43*(1), e12702. https://doi.org/10.1111/cogs.12702

Syrett, K., LaTourrette, A., Ferguson, B., & Waxman, S. R. (2019). Crying helps, but being sad doesn't: Infants constrain nominal reference online using known verbs, but not known adjectives. *Cognition*, *193*, 104033. https://doi.org/10.1016/j.cognition.2019.104033

Tovar-Perez, P., Smith, G., Alarcon Chiroboga, M., Gaston-Panthaki, A., Castillo, A., María Fernández, A., & Childers, J. B. (2025). How often do verbs co-occur with relevant events? Examining parent and toddler speech in latin america and the us. *Journal of Cognition and Development*, *26*(2), 198–220. https://doi.org/10.1080/15248372.2024.2418628

Yu, C., Zhang, Y., Slone, L. K., & Smith, L. B. (2021). The infant's view redefines the problem of referential

uncertainty in early word learning. *Proceedings of the National Academy of Sciences*, *118*(52), e2107019118. https://doi.org/10.1073/pnas.2107019118

Yuan, S., & Fisher, C. (2009). "Really? She blicked the baby? ": two-year-olds learn combinatorial facts about verbs by listening. *Psychological Science*, *20*(5), 619–626. https://doi.org/10.1111/j.1467-9280.2009.02341.x