

The Role of Syntactic and Referential Evidence in Verb Learning across Exposures

Yiran Chen (chen2537@wisc.edu)

Waisman Center, University of Wisconsin-Madison,
1500 Highland Ave, Madison, WI 53703 USA

Alexander LaTourrette (alatourret@haverford.edu)

Department of Psychology, Haverford College
370 Lancaster Avenue, Haverford, PA 19041

John Trueswell (trueswel@psych.upenn.edu)

Department of Psychology, University of Pennsylvania,
425 S University Ave, Philadelphia, PA 19139 USA

Abstract

Early word-learning opportunities are often highly ambiguous, with this problem being especially difficult for verbs. While a verb's syntax can help to identify the referent event from the environment, learners still need to contend with temporal and spatial misalignment between verbs and their referent events. Although children are shown to use syntax to infer verb meaning when there is initially no co-occurring referent event, it remains unclear what role syntax plays in verb learning across exposures in tandem with referential information. With three adult word-learning experiments, we showed that while syntax independently informed verb meaning in the absence of referents, it did not additionally constrain subsequent mappings when a referent was present. These results reveal both the power of syntax in cross-situational verb-learning—persisting across exposures—and its limitations—failing to supersede co-present referents.

Keywords: Verb learning, Syntactic bootstrapping, Cross-situational learning, Acquisition, Psycholinguistics

Introduction

People only sometimes talk about visually co-present events (e.g., “She’s running!”). Instead, we often talk about events from the past (“She left yesterday”) or in the future (“Let’s tie our shoes”). This makes learning verb meanings particularly challenging: we only sometimes hear a verb in the presence of a perceptually salient referent (Gillette et al., 1999; Medina et al., 2011; Gleitman, 1990). And even when word and world temporally align, a scene can be described from different perspectives (“He is giving him the ball” vs. “He is getting the ball from him”) – making verbs hard to learn from observation alone (Fisher et al., 1994).

Fortunately, learners can rely on additional sources of information to infer verb meaning. One important piece of evidence is the linguistic context of the verb. A now robust literature reveals that the syntactic environment of a verb may help verb learning by acting as a “zoom lens,” directing learners’ attention to the event with matching structure (Gleitman, 1990; Landau & Gleitman, 1985; Naigles, 1990, 1993; Nappa et al., 2009, Gleitman et al., 2005, Fisher et al., 2010). For instance, verbs which refer to causal events are typically used in transitive frames (e.g., “She pushed him”)

while verbs referring to events with non-causal events are typically used in intransitive frames (e.g., “She smiled”). Thus, by using the sentence frames that a new word occurs in, learners can gain insight into the word’s meaning — a process known as “syntactic bootstrapping” (Gleitman, 1990; Landau & Gleitman, 1985; Naigles, 1990, 1993).

While syntax is helpful to verb learning in-the-moment, it remains less clear how syntax might facilitate word learning across several occurrences of a word. One prediction of the syntactic bootstrapping theory is that learners can use the multiple frames in which a verb occurs to infer a verb’s semantic subclass (Fisher et al., 1991; Landau & Gleitman, 1985; Gleitman, 1990). For example, the distributional profile of the verb “see,” which takes both NP and CP complements, may signal to blind children that “see” is a perception verb (Landau & Gleitman, 1985). It is a crucial question, then, how children track and synthesize syntactic information across exposures in order to successfully update a verb’s meaning in new learning instances.

Importantly, recent work has shown that children can retain syntactic constraints on a word’s meaning across delays, suggesting syntax may play a role across instances. In Yuan and Fisher (2009), two-year-olds heard a novel verb used in a transitive frame during a dialogue between two women without any co-present event referents. Toddlers were then more likely to map that verb to a two-participant causative event on a later exposure, compared to toddlers who heard the verb used in an intransitive context during the dialogue (see also Arunachalam & Waxman, 2010).

Current Work

Though these findings showcased the power of syntax for verb inference after a delay, it remains unknown how inferences from syntax interact with a verb’s referential context across occurrences (Yuan & Fisher, 2009). While children will sometimes hear a verb in the absence of referents and have to guess its meaning based on linguistic cues alone, they often also encounter referential evidence that could be used and retained across occurrences.

It is crucial, then, to ask what role syntax plays in verb learning across different referential contexts: that is, whether inferences from linguistic information can be used to infer meanings across different, referentially rich learning instances. For instance, if a learner observed a playground rife with different activities, and heard “Look, he’s daxing her!”, this might lead our learner to several inferences: she might note that “dax” is used in a transitive sentence; she might infer that “dax” refers to a causative event; she might even map “dax” to a specific referent event (e.g., a pushing event). But if the learner next encounters “dax” in a different context and realizes her previous hypothesized meaning (“push”) was incorrect, little work has examined which, if any, of the above inferences she might use to successfully generate a new guess.

To address this, the current study examines syntax’s role in verb learning across multiple referential contexts with a series of experiments with adults as a first step. Adults reliably outperform children in cross-situational word learning tasks (Benitez & Li, 2023; Fitneva & Christiansen, 2015), and they are also syntactically sophisticated and equipped with better memory than children (e.g., Vlach & Sandhofer, 2012). Thus, we expected adults to present the best-case scenario for detecting syntactic bootstrapping effects across occurrences of a verb. Adult performance itself is also of interest, given that adult second-language learners face the same learning and computational problem: what information should be tracked, stored, and synthesized across verb occurrences? Thus, findings may also inform adult second language acquisition (see Hulstijn 2003 for a review).

Across three experiments, we exposed adults to a series of novel verbs, each presented in either a transitive frame (“The girl daxes the boy”) or a conjoined-subject intransitive frame using the adverbial modifier “together” (“The boy and the girl dax together”). Prior work has shown that both frames are excellent cues to a verb’s meaning. As noted above, transitive frames reliably lead learners to map verbs to causative events (e.g., “kicking”), even after a delay (Yuan & Fisher, 2009; Arunachalam & Waxman, 2010). Both adults and children also show a weaker, but still substantial, preference for mapping verbs presented in the conjoined-subject intransitive frame to non-causative events (e.g., “waving”) (Arunachalam et al., 2016).

In each study, our crucial test was whether learners would leverage those initial, syntactically informative frames to infer the verb’s causal structure. If so, this should guide inferences not just on the initial exposure but on a subsequent exposure as well—when any initially available referents were no longer present. If participants use previous syntactic cues to inform their subsequent mappings, their selections on this second exposure should match the causal structure suggested by the preceding syntactic frame.

With this as our guiding prediction, we conducted three experiments examining different learning scenarios. In

Experiment 1, we asked if learners retain syntactic inferences even when the initial referent becomes unavailable, comparing this against an absent-referent scenario. In Experiment 2, we then tested the contribution of syntax to these inferences, asking whether an informative frame, compared to an uninformative one, facilitates verb learning when all learners see the same initial referent. Finally, Experiment 3 tested whether when a verb’s syntax mismatches the salient referential event in view, learners rely more on the syntactic constraints or the co-occurring events in subsequent inferences about the verb’s meaning.

Experiment 1

In Experiment 1, we tested whether adult learners retain syntactic inferences when the verb’s initial referent becomes unavailable, comparing this scenario against one in which there is no referent initially available. To illustrate the importance of such inferences, consider the following scenario. Suppose a learner hears the verb “push” for the first time embedded in the sentence “He is pushing her.” Based on prior literature, we expect a learner would map “push” to a causative event occurring in view. However, unlike most studies, learners’ daily lives offer no guarantees that the visible causative event will be the correct mapping. For example, perhaps there is a salient carrying event in view instead (e.g., an outnumbered parent who is pushing one child in the stroller while carrying another). Learners might then reasonably, if incorrectly, identify carrying as the meaning of “push.” However, all is not necessarily lost: in addition to remembering the incorrect “push”-carry mapping, learners might remember that “push” occurred in a transitive frame, or that “push” referred to a causative action. Then, the next time the learner hears “Look, pushing”, and no carrying event is present, perhaps our learner would be able to use the prior exposure’s syntactic constraints to map “push” to a new, co-occurring causative event, rather than a non-causative event. In this way, learners could leverage syntactic information from prior exposures – even if their initial meaning proved incorrect. Experiment 1 tested exactly this question, using a novel verb learning paradigm.

Methods

Participants Eighty native speakers of American English recruited from the University of Pennsylvania subject pool participated for course credit.

Materials Stimuli were video clips depicting novel actions that served as potential referents for novel verbs and audio clips of English sentences that contained the novel verbs.

We created 36 short video clips, each of which depicted a novel action and lasted 1-3 seconds. One actor and one actress were present in every video. Half the clips depicted novel causative actions, and the other half non-causative actions. In causative events, one actor caused the other to

take an action (e.g., the boy crosses the girl’s arms). In non-causative events, both actors performed the action synchronously (e.g., both synchronously do the superman pose). The causer’s identity was counterbalanced (Fig. 1).

The English sentences were recorded by a female native speaker of American English in infant-directed speech¹. The audio clips had an average duration of 10 seconds. There were two types of frames – transitive (e.g., “Sometimes, the boy lorp the girl. The boy really likes to lorp the girl!”) or conjoined-subject intransitive (“Sometimes, the boy and the girl fep together. They really like to fep together!”).

Procedure All experiments were conducted online using PCIbex (Zehr & Schwarz, 2018). The experiment consisted of three blocks identical in structure. Each block featured 4 distinct novel verbs. Thus, adults encountered 12 verbs in total. Participants heard each verb first in an Exposure trial and then again in a Test trial. During the Exposure trial, half the verbs were presented in transitive sentences (red in Figure 1) and half in conjoined-subject intransitive sentences with the “together” modifier (blue in Figure 1).




Figure 1. One sample exposure block in Experiment 1.

This syntactic manipulation was crossed with our referential manipulation: during Exposure, half of the trials had no referent event present (No-referent Condition) and the other half had referent events that matched the syntax - i.e., causative events shown with transitive frames and non-causative events with intransitive frames (Match Condition). This manipulation was done within-subjects to create a more sensitive test and reflect the fact that everyday language environments feature multiple frame types and a mix of present and absent referents. On trials with referent events, the video referent events were on loop until the audio finished playing. On trials with no referents, participants only saw a closed curtain, which they had been told indicated they could not see the event the sentence was describing (Figure 1). The four trials in each block were arranged in a pseudo-random order, and verbs’ syntactic-frame and reference-condition assignments were counterbalanced across experimental lists.

After one block of exposure trials, participants were tested on the four verbs they just learned. For each verb, participants were prompted with the verb in a gerund frame

(i.e. “Where’s fepping? Find fepping!”) and given the chance to map it to either a novel causative or a novel non-causative event (Figure 2). Note that participants had seen neither of the test events before. Participants had unlimited time to select one as the verb’s referent. To avoid specific association between certain event pairs, which events appeared at exposure or test as well as which events were paired with the same verb were counterbalanced via experimental lists. If participants use a verb’s previous syntactic frame to infer its meaning even when the original referent is unavailable, participants’ selections should match the causal structure suggested by previous syntax, even though the event itself is novel.



Figure 2. A sample test trial.

Analysis To test whether learners in each condition showed a preference for the novel event that matched the causativity suggested by previous syntax, we built a logistic mixed-effect model to predict causativity match (0 vs 1). We included Condition (No-referent vs Match), Frame type (Transitive vs Intransitive) and the interaction as deviation-coded fixed effects. Following Barr et al., (2013), we included the maximal random structure: by-subject random intercepts and random slopes of Condition and Frame type and by-item random intercepts and by-item random slopes of Condition and Frame type.

Results

First, this model revealed a significant positive intercept ($\beta=0.652$, $SE =0.165$, $p<0.001$), suggesting that learners’ selection at Test conformed to the causativity suggested by the verb’s syntactic frames on the initial Exposure trial. That is, adults preferred to map verbs previously heard in transitive frames to causative events, and verbs previously heard in intransitive frames to non-causative events.

Notably, we did observe a more reliable causativity-matching effect for verbs previously heard in transitive frames than verbs previously heard in intransitives ($\beta=-0.202$, $SE=0.094$, $p=0.031$). This suggests learners’ preference for mapping non-causative events to verbs used in conjoined-subject intransitive frames with “together” as a modifier (cf. Arunachalam et al., 2016) is not as robustly retained across exposures. This is also consistent with prior research that shows that transitive sentences are a more informative cue to event structure than (unmodified) intransitive sentences (Arunachalam & Waxman, 2010).

However, as shown in Figure 3, participants’ performance did not differ across the two referential conditions ($\beta=0.059$, $SE=0.099$, $p=0.554$, $M_{Match}=64.6\%$, $M_{No-referent}=61.2\%$).

¹ This was in preparation for a future study with toddlers.

There was also no significant interaction between Condition and Frame-type ($\beta=0.143$, $SE=0.119$, $p=0.232$). Notably, the No-referent Condition's above-chance performance converges well with previous findings that, in the absence of a referent, syntax informs subsequent verb mappings (Yuan & Fisher, 2009). The Match Condition's performance, on the other hand, provides new evidence that adults make new semantic inferences in line with previous syntax, even when the original referent is unavailable on the later exposure.

Taken together, these findings show that syntax guides verb mapping on a subsequent exposure both when no referent is present on the first exposure and when a referent is present at exposure but does not reappear later. However, when a referent is available, there are at least two possible inferences that might yield syntax-compatible referent selections on subsequent exposures: learners could rely on either the previous syntax (e.g., remembering that "dax" occurred in a transitive frame) or on the causativity of the referent (e.g., remembering "dax" referred to a causative event). Experiment 2 tested whether learners rely on the previous syntax or only on the previous referent mapping.

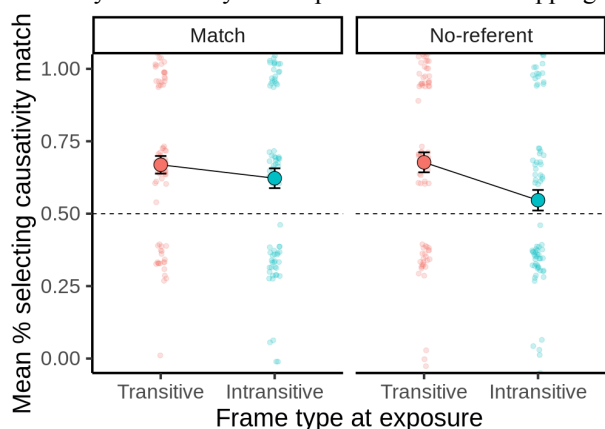


Figure 3. Mean proportion of selecting the causativity match event (i.e., causative events for verbs heard in transitive trials and non-causative events for verbs heard in intransitive trials) at test in Expt. 1. Dashed line indicates chance selection (50%). Error bars represent SEM.

Experiment 2

To test the role of syntax in learning verbs across exposures, Experiment 2 modified Expt. 1's paradigm. Now, we always presented a referent event at Exposure but varied whether the syntax paired with the event provided informative cues to the verb's causal structure (e.g., "The boy daxes the girl") or no informative cues (e.g., "Look, there's daxing"). If syntax facilitates later verb inference, learners should make more causative matches after informative frames.

Methods

Participant Eighty native speakers of American English recruited from the University of Pennsylvania subject pool participated for course credit.

Procedure Expt. 2 adopted a similar design to Expt. 1 with two changes. First, while participants still learned and were tested on 12 novel verbs across three blocks, this time, during Exposure trials, a referent event was always present. Half of the verbs were paired with causative actions during exposure, the other half with non-causative actions.

Second, for each learner, half the verbs were presented with informative, matched syntax at Exposure (as in Expt. 1's Match Condition) whereas the other half were presented with Under-informative syntax, (i.e., a gerund frame) (Figure 4). The Under-informative frame still mentioned the agents (i.e. "The boy and the girl") so that subject NPs would be matched across conditions (cf. Arunachalam & Waxman, 2010; Nobel et al., 2011). Verb conditions were counterbalanced across experiment lists. Test trials were identical to Expt 1 (Figure 2): we tested if learners would select a new referent event for the verb that matched the Exposure referent's causativity and, in the Match Condition, also the syntactic frame heard at Exposure. If learners use previous syntax, they should make more causative-matches at Test in the Match than the Under-informative condition.




Figure 4. A sample transitive and intransitive trial as presented in the Match or Under-informative conditions.

Analysis To test whether learners in each condition showed a preference for the novel event that matched the causativity of previous referent, we built a logistic mixed-effect model to predict causativity match (0 vs 1) with Condition (Under-informative vs Match), Referent Causativity (Causative vs Non-causative) and their interaction as deviation-coded fixed effects. Random structure included a by-subject random intercept and random slopes of Condition and Referent Causativity and by-item random intercepts and random slopes of Condition and Referent Causativity.

Results

Overall, learners were likely to select the novel referents at Test that matched the causativity of the verb's referent at Exposure: we observed a significant positive intercept in the model ($\beta=0.502$, $SE=0.111$, $p<0.001$, See Figure 5). This

model also revealed a significant effect of Referent Causativity ($\beta=0.229$, $SE=0.091$, $p=0.012$), indicating learners were more likely to map a verb paired with a causative event at Exposure to a new causative event at Test than to map a verb paired with a non-causative event at Exposure to a new non-causative event at test.

However, the model found no significant difference in performance in the two conditions ($\beta=0.089$, $SE=0.081$, $p=0.268$, $M_{Match}=59.1\%$, $M_{Under-informative}=62.6\%$). There was also no significant interaction between Condition and Referent Causativity ($\beta=0.025$, $SE=0.075$, $p=0.744$). This suggests that the informative syntax didn't have any additional impact on learners' referent selections at Test.

In summary, the Match Condition replicated that of Expt. 1, providing robust evidence that even when participants' original word-referent mapping is no longer available, they continue to select new referents that are compatible with the prior exposure's syntax and the causativity of the original referent. However, the Under-informative condition's similar performance revealed that this causativity-matching effect on subsequent verb inferences may be driven by the prior exposure's referent, rather than by the prior exposure's syntax. Learners mapped verbs to causativity-matching referents even without informative syntax at Exposure. Overall, in the presence of a referent, syntax didn't additively influence learners' subsequent verb mapping.

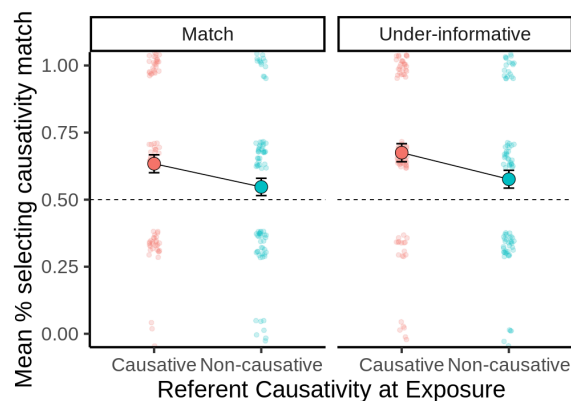


Figure 5. Mean proportion of selecting the causativity match event (i.e., the causative events for transitive trials and non-causative events for intransitive trials) at test in Expt. 2. Dashed line indicates chance probability (50%). Error bars represent standard errors of the mean.

Experiment 3

Experiments 1 and 2 provided evidence that both syntactic cues (Expt. 1) and referential cues (Expt. 2) can guide subsequent verb mappings. In Experiment 3, we compared the strength of these cues, asking what happens if the event that is salient in learners' visual world when a verb is uttered does not match the syntax they hear: would learners still rely on the syntax in this situation? Or would they instead rely on the referent alone and discard the syntactic cue, as suggested by Expt. 2?

Methods

Participant Eighty native speakers of American English recruited from the University of Pennsylvania subject pool participated for course credit.

Procedure Expt. 3 is identical to Expts. 1 and 2 with one critical change. For each learner, on half of the Exposure trials, the syntax matched the causal structure of the referent (as in the prior studies' Match Condition), but on the other half of trials, in the Mismatch Condition, the syntax and referent were mismatched (i.e. intransitive sentences accompanied causative actions, and transitive sentences accompanied non-causative actions) (Figure 6). For each verb, half of participants learned it in the Match Condition and the other half learned it in the Mismatch Condition. The test was identical to the previous two experiments (Figure 2): participants were asked to map each learned verb to either a novel causative or non-causative action.

Participants' selection at test was analyzed in the same way as in Experiment 2, with a mixed-effects model predicting whether learners selected a referent at Test that matched the causativity of the Exposure trial's referent (note that in the Mismatch Condition, if learners are guided by syntax, they should show a preference for the Test referent which does *not* match the Exposure referent in causativity).

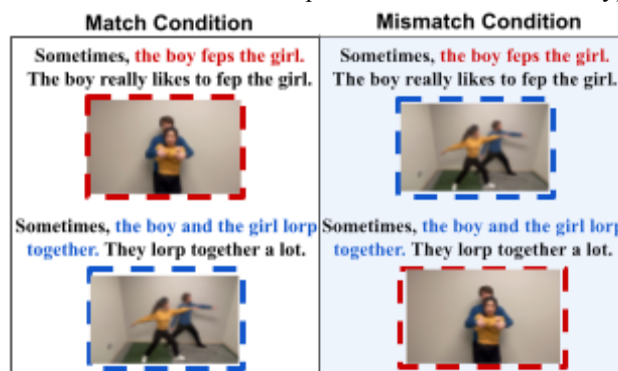


Figure 6. One sample transitive and intransitive trial in the Match and Mismatch Conditions.

Results

Experiment 3 asked whether learners given mismatched cues from syntax and referent at exposure would rely more on the syntax or the referent on a subsequent exposure. Our results (see Figure 7) suggest learners made inferences in line with the preceding referent, not the preceding syntax. A significant positive intercept in the model suggests that learners were likely to select the novel actions that were matched in causativity with the referent at Exposure: ($\beta=0.625$, $SE=0.125$, $p<0.001$). In line with the previous two experiments, learners also showed stronger causativity matching for verbs previously paired with causative events compared to verbs previously paired with non-causative events ($\beta=0.202$, $SE=0.095$, $p=0.033$).

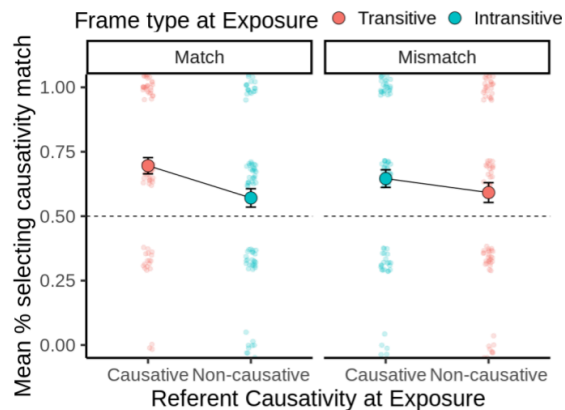


Figure 7. Expt. 3 participants' mean proportion of selecting the event at test matching the exposure event's causativity (e.g., for a verb paired with a causative event at exposure, choosing a causative event at test). Dashed line indicates chance probability (50%). Error bars represent SEM.

More critically, we observed no significant effect of condition ($\beta=0.030$, $SE=0.122$, $p=0.806$, $M_{Match}=63.3\%$, $M_{Mismatch}=61.9\%$). That is, learners were no more likely to choose the event matching the prior referent's causativity when they received matched versus mismatched syntax at Exposure. Indeed, a mixed-effects analysis of the Mismatch Condition alone revealed that even in this condition, participants significantly preferred the test event matching the prior referent's causativity over the event matching the prior syntax ($\beta=0.582$, $SE=0.143$, $p<0.001$). Finally, no significant interaction emerged between Condition and Referent Causativity ($\beta=0.012$, $SE=0.099$, $p=0.900$). Across conditions, then, learners largely did not rely on the syntactic cues presented during Exposure when selecting a new referent at Test. Instead, they adhered to the causal structure of the verb's referent at Exposure.

General Discussion

Across three experiments, adults successfully leveraged information from a verb's previous exposure to make a new guess about its meaning. In Experiment 1, adults successfully used syntactic cues presented at Exposure in the absence of a referent to identify syntactically compatible referents on the next exposure (No-referent Condition). Moreover, when a referent was present at Exposure, we found that adults mapped verbs to new events that matched that previous exposure's referent in causativity. Adults made these referent-based mappings to a similar extent regardless of whether the previous exposure had included informative syntactic cues (All Expts: Match Condition), uninformative syntactic cues (Expt. 2: Under-informative Condition), or mismatching syntactic cues (Expt. 3: Mismatch Condition). Thus, learners robustly attend to the causativity of a verb's referent event and, in the absence of such an event, to the causativity suggested by a verb's syntactic frame. Learners

then generate new meanings consistent with that causativity on subsequent trials. Thus, causativity is a salient dimension in both linguistic and referential domains.

However, these studies show mixed results on the impact of syntactic cues across learning instances. Converging with prior work (e.g., Arunachalam & Waxman, 2010; Yuan & Fisher, 2009), the results of Experiment 1 indicate that syntax can guide subsequent verb mappings when no referent is present on the exposure. However, even in this condition, transitive frames showed a much more robust impact on subsequent verb mappings than conjoined-subject intransitive frames modified by "together." Prior work has found that both types of frames are reliably used by learners within a single exposure to map verbs to, respectively, causative or non-causative actions (cf. Arunachalam et al., 2016). Thus, the current results suggest there may be variation in how well different syntactic inferences are retained across exposures—even in the absence of a referent.

Moreover, while the results of Expt. 1 indicate that syntax (or at least, transitive frames) can guide subsequent verb mappings when no referent was present, subsequent studies indicated that when a referent event is present at exposure, syntax does not reliably affect subsequent mappings. Specifically, syntax had no additive effect when a syntactically compatible referent event was present (Experiment 2), and it did not have an inhibitive effect when a mismatched event was present (Experiment 3).

While these results are very much only a first step, they raise new questions about the role of syntax in verb learning across exposures. Syntax is an undeniably important resource for word learning, helping learners to "zoom in" on viable candidate meanings when a verb is uttered (e.g., Naigles, 1990). Importantly, the current results also suggest that even when these initial, syntactically informed mappings are incorrect, learners still use the mappings (or at least, their causativity) as a guide to generating new semantic hypotheses on subsequent exposures. Thus, even if the syntax itself is not retained across exposures, it may continue to indirectly guide subsequent verb learning.

That said, it remains possible that in some cases, a verb's syntactic history is used to infer the verb's meaning on a subsequent exposure: in particular, our results suggest this might occur primarily when no referent is present. In other cases, however, learners may take a salient referent as a more important cue than a verb's syntax and rely on it whenever a referent is available.

In sum, these results reveal both the power of syntax in verb-learning – persisting across exposures without referents – and its limitations – failing to supersede referents when they are present. While these findings with adult learners directly inform second language acquisition, it is possible children may employ different strategies when presented with these conflicting cues. In future work, we plan to use a similar paradigm to test how children use syntax across exposures in first language acquisition.

References

- Arunachalam, S., & Waxman, S. R. (2010). Meaning from syntax: Evidence from 2-year-olds. *Cognition*, 114(3), 442-446.
- Arunachalam, S., Syrett, K., & Chen, Y. (2016). Lexical disambiguation in verb learning: Evidence from the conjoined-subject intransitive frame in English and Mandarin Chinese. *Frontiers in psychology*, 7, 138.
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of memory and language*, 68(3), 255-278.
- Benitez, V. L., & Li, Y. (2023). Cross-situational word learning in children and adults: The case of lexical overlap. *Language Learning and Development*, 1-24.
- Fisher, C., Gleitman, H., & Gleitman, L. R. (1991). On the semantic content of subcategorization frames. *Cognitive psychology*, 23(3), 331-392.
- Fisher, C., Hall, D. G., Rakowitz, S., & Gleitman, L. (1994). When it is better to receive than to give: Syntactic and conceptual constraints on vocabulary growth. *Lingua*, 92, 333-375.
- Fisher, C., Gertner, Y., Scott, R. M., & Yuan, S. (2010). Syntactic bootstrapping. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(2), 143-149.
- Fitneva, S. A., & Christiansen, M. H. (2017). Developmental changes in cross-situational word learning: The inverse effect of initial accuracy. *Cognitive Science*, 41, 141-161.
- Gillette, J., Gleitman, H., Gleitman, L., & Lederer, A. (1999). Human simulations of vocabulary learning. *Cognition*, 73(2), 135-176.
- Gleitman, L. (1990). The structural sources of verb meanings. *Language Acquisition*, 1(1), 3-55.
- Gleitman, L. R., Cassidy, K., Nappa, R., Papafragou, A., & Trueswell, J. C. (2005). Hard words. *Language Learning and Development*, 1(1), 23-64.
- Hulstijn, J. H. (2003). Incidental and intentional learning. *The handbook of second language acquisition*, 349-381.
- Landau, B., & Gleitman, L. R. (1985). *Language and experience: Evidence from the blind child*. Cambridge, MA: Harvard University Press.
- 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.
- Naigles, L. (1990). Children use syntax to learn verb meanings. *Journal of Child Language*, 17(2), 357-374.
- Naigles, L. G., & Kako, E. T. (1993). First contact in verb acquisition: Defining a role for syntax. *Child Development*, 64(6), 1665-1687.
- Nappa, R., Wessel, A., McEldoon, K. L., Gleitman, L. R., & Trueswell, J. C. (2009). Use of speaker's gaze and syntax in verb learning. *Language Learning & Development*, 5(4), 203-234.
- Noble, C. H., Rowland, C. F., & Pine, J. M. (2011). Comprehension of argument structure and semantic roles: Evidence from English-learning children and the forced-choice pointing paradigm. *Cognitive Science*, 35(5), 963-982.
- Vlach, H., & Sandhofer, C. M. (2012). Fast mapping across time: Memory processes support children's retention of learned words. *Frontiers in psychology*, 3, 17291.
- 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.
- Zehr, J., & Schwarz, F. (2018). PennController for Internet Based Experiments (IBEX).