

Specifying the Role of Linguistic Information in Verb Learning

Sudha Arunachalam and Sandra R. Waxman
Northwestern University

Decades of research have documented that acquiring the meanings of verbs is, on average, more difficult than acquiring the meanings of nouns. This is likely because verb learning requires more information, and information of a different kind, than noun learning. One particularly powerful source of information for verb learning is linguistic context. For example, Naigles (1990) demonstrated in a classic study that 2-year-olds infer that verbs appearing in transitive sentences carry causative meanings. The present study explores in more detail what aspects of the linguistic information that is typically carried in a transitive sentence—the syntactic structure itself, and/or the semantic content of the noun phrases introducing the arguments—are most useful for these young verb learners.

1. Information sources for identifying verb meaning

A clever research paradigm using adults in a simulated word learning task has offered insight into the kinds of information that are particularly potent for identifying the meaning of different kinds of words (Gillette et al., 1999; Snedeker & Gleitman, 2004). These studies investigated the conditions under which adults could determine the referent of a word, given varying amounts of information. Adult observers viewed video clips of naturalistic interactions between a parent and infant. The video clips were silent, except for a single beep, during which the adult participants were asked to guess what word the parent had uttered to the infant at that moment. Participants' identification of the words was much more successful for nouns than for verbs, providing support for the hypothesis that acquiring a word's meaning from observation alone is more difficult for verbs than for nouns. The fact that this 'noun advantage' was present in adults, who are fully capable of representing the concepts underlying verb meaning, documented that the obstacles that young learners encounter in verb learning cannot be accounted for by appealing simply to the inherent difficulties representing verb concepts; instead, the challenges are also intrinsically tied to the informational requirements of verb learning.

These studies further documented that linguistic information is a particularly informative cue for identifying verbs in this task, and also went further to offer insight into the particular *kinds* of linguistic information that are

* This research was funded by NIH HD30410 to the second author.

most potent. In some conditions, in addition to observing the visual scenes, adult participants were given linguistic information to aid them in guessing the word. For some participants, this linguistic information consisted of a list of the nouns that had occurred in the sentence containing the verb. For others, the linguistic information was the syntactic frame in which the verb appeared: the content words were replaced, Jabberwocky-style, with nonsense words, so that the frame only conveyed syntactic information. For a third group of participants, the linguistic information included both of these information sources. The amount and type of linguistic information provided had a clear effect: Participants hearing the list of nouns guessed the verbs 29% of the time; those hearing the syntactic frame performed significantly better, guessing the verbs 53% of the time; and those given both the nouns and the syntax performed best, guessing the verbs 78% of the time (Snedeker & Gleitman, 2004). These results demonstrate that linguistic information is instrumental in determining a verb's referent, and further, that different kinds of linguistic information (syntactic frame, argument labels) can be differentially potent. For adults, both cues in combination were most informative.

This simulated task was, of course, designed to depict the best possible learning scenario, in which adults are not required to *learn* anything, but rather to guess words for which they have already established meanings. Evidence from 3- and 7-year-olds in a very similar simulated word guessing paradigm provides a first indication that for children, too, linguistic information is important for guessing verbs (Piccin & Waxman, 2007; Waxman & Plotkin, 2008). These simulation tasks thus provide an *in principle* demonstration of the importance of linguistic information, and particular kinds of linguistic information, for verb learning.

But *in practice*, young learners encountering a new word for the first time in the context of an observational scene face the additional task of determining how words map to the world. Although we know that toddlers can use linguistic information to acquire verb meanings, it may nonetheless be challenging for them to coordinate their working memory and processing capacities effectively and efficiently enough to use the linguistic information they are given. It is therefore important to determine how toddlers use different linguistic cues independently and in concert.

2. Verb learning in toddlers

The current study inquires into two types of linguistic information—syntactic frame and semantic content of argument labels—that have been implicated in previous work with adults and older children as valuable in discovering verb meaning. We use a novel word learning task, a well-established method for studying vocabulary acquisition in young children, to explore the potency these two types of linguistic information.

In recent work, Waxman, Lidz, Braun and Lavin (2009) established a paradigm that elicits successful verb learning in 2-year-olds, demonstrating that

toddlers can use the linguistic context in which a novel word is presented to determine its grammatical category (either a noun or verb) and map it to an appropriate referent (Waxman et al., 2009). Arunachalam and Waxman (2009) used a very similar paradigm to look more closely at the sources of information that were essential to toddlers' success. See Table 1. Toddlers were first familiarized to either novel nouns or novel verbs in dialogue scenes. The dialogues allowed toddlers time to hear the novel words and determine their grammatical categories from the linguistic context, before viewing a relevant event (see Yuan and Fisher, 2009). Next, toddlers were presented with several visual scenes in which an actor performed a simple causative action on an object (e.g., waving a balloon). The linguistic information toddlers heard in conjunction with these visual scenes varied by condition: Those in the Noun condition heard, "Look! The boy is waving a *pilker!*", those in the Verb condition heard "Look! The boy is *pilking* a balloon!" Toddlers then viewed a contrast phase, in which the actor was performing a very different action with a very different object, to alert them that not all possible scenes were candidate referents of the novel words.







Dialogue	Familiarization		Test	
	Multiple Exemplars	Contrast		
				 Familiar Object  Familiar Action
<p>Noun: Hey! The girl painted the <i>pilker!</i></p> <p>Verb: Hey! The girl <i>pilkerd</i> the shirt!</p>	<p>Full Noun Phrases Noun: Look, the boy is waving a <i>pilker!</i></p> <p>Verb: Look, the boy is <i>pilking</i> the balloon!</p> <p>Pronouns Noun: Look, he's waving a <i>pilker!</i></p> <p>Verb: Look, he's <i>pilking</i> it!</p>	<p>Noun: Uh-oh, that's not a <i>pilker.</i></p> <p>Verb: Uh-oh, he's not <i>pilking</i> that.</p>	<p>Noun: Yay, that is a <i>pilker!</i></p> <p>Verb: Yay, he is <i>pilking</i> that!</p>	<p>Noun: Where's the <i>pilker?</i></p> <p>Verb: Where's <i>pilking?</i></p>

Table 1. Representative stimuli from Arunachalam & Waxman (2009)

At test, all toddlers saw two scenes side-by-side. In one, the actor was performing the now-familiar action, but on a new object (e.g., waving a rake); in the other the actor was performing a new action, but on the familiar object (e.g., tapping the balloon). Note that to succeed at this task, toddlers in the Verb condition had to accept an event that preserved the action, despite a change in objects, and those in the Noun condition had to accept an event that preserved the object, despite a change in the action in which it was involved. Toddlers succeeded, demonstrating that their initial representations of the meanings of these novel nouns and verbs were sufficiently abstract to permit them to map the novel words and extend them appropriately.

However, the particular linguistic context in which novel verbs were presented determined toddlers' success in the task. When the novel words appeared in sparse semantic contexts, in which the full noun phrase arguments

were replaced with pronouns, toddlers still successfully mapped novel nouns, but now failed to map novel verbs. In the Noun condition, toddlers heard “Look! He’s waving a *pilker!*”, and in the Verb condition, they heard “Look! He’s *pilking* it!” This suggests that the semantic content inherent in the noun phrases labeling the verb’s arguments (e.g., *the balloon* vs. *it*) helps toddlers discover verb meaning, over and above the benefits of informative syntactic contexts.

In the current study, we pursue this finding in two ways. First, we systematically uncouple the syntactic and semantic information to identify the contribution of each source of linguistic information—syntactic frames and semantically rich noun phrases—in verb learning. And second, we ask whether toddlers require the multiple exposures to the novel verbs and scenes that we provided in our earlier experimental paradigm (Arunachalam & Waxman, 2009).

3. Experiment 1: Manipulating Syntactic Richness

In Experiment 1, we focused on the contribution of syntactic frame. All toddlers were given semantically rich noun phrases labeling the verb’s arguments; what varied was whether they were presented in an informative syntactic context. In the [+Nouns –Syntax] condition, the nouns appeared outside the sentence introducing the verb: “Let’s see a boy, and a balloon. Let’s see *pilking!*” In the [+Nouns +Syntax] condition, they appeared embedded in a transitive sentence: “A boy is gonna *pilk* a balloon. Let’s see!” See Table 2.

The stimuli and procedure were very similar to Arunachalam & Waxman (2009), except that we drastically reduced the amount of exposure toddlers were given to the novel verbs and scenes. Toddlers heard the novel verbs just twice during familiarization (only once in an informative linguistic context), and saw only a single instance of the event described by the verb, viewed twice on either side of the screen. See Table 2.

3.1 Methods

Participants. Thirty-six typically-developing toddlers (19 males) with a mean age of 26.9 months (range: 25.0-29.9) were included in the final sample. All were recruited from Evanston, IL and surrounding communities. They were acquiring English as their native language and spent less than 25% of the time hearing another language. Caretakers completed the MacArthur Long Form Vocabulary Checklist: Words and Sentences (Fenson et al., 1993). Toddlers’ mean production vocabulary was 507 words (range: 140 to 679); there were no differences in vocabulary between the conditions. To insure that toddlers included in our final sample were capable of providing systematic pointing responses, we developed stringent inclusion criteria: toddlers had to point correctly on at least two of four training trials. 11 toddlers failed to meet these criteria and were replaced. Two additional toddlers were excluded due to fussiness and one due to experimenter error.

Materials. The materials were adapted from Waxman et al. (2009). The visual materials were digitized recordings of human actors performing continuous actions on inanimate objects. These were edited to create the sequences described in Table 1 and were presented on a 20 in. screen.





	Linguistic Familiarization	Event Familiarization	Test	
Exps 1 & 2	 (still image) Boy with balloon	 (dynamic scene) Boy waving balloon	 Boy tapping balloon <i>(Familiar Object)</i>	 Boy waving rake <i>(Familiar Action)</i>
Exp 1	<u>+Nouns –Syntax</u> Let's see a boy, and a balloon. Let's see <i>pilking</i> . <u>+Nouns +Syntax</u> A boy is gonna <i>pilk</i> a balloon. Let's see.	Look, <i>pilking</i> !	Now look, they're different. Where's <i>pilking</i> ?	
Exp 2	<u>–Nouns +Syntax</u> Let's see what happens now. He's gonna <i>pilk</i> it. <u>+Nouns +Syntax</u> Let's see a boy, and a balloon. He's gonna <i>pilk</i> it.			

Table 2. Representative set of the stimuli presented in Experiments 1 and 2

The linguistic materials were recorded by a female native English speaker in a child-directed speech register. These utterances were edited to control duration and amplitude, and were synchronized with the visual stimuli. The auditory stimuli were presented via a speaker directly below the center of the screen.

Apparatus and procedure. Toddlers and their caretakers were welcomed into a playroom. While the toddler and experimenters played with toys, the caretaker signed a consent form and completed the vocabulary checklist. The toddler and caretaker then accompanied the experimenters into an adjoining test room. The toddler was seated in an infant seat, 12 in. from the screen. The caretaker sat behind the toddler, and was asked not to talk during the study or otherwise influence their toddler's behavior. One experimenter controlled the experimental procedure from behind a curtain, while the other, seated beside the toddler, elicited pointing responses. Toddlers' pointing behavior was recorded with a video camera centered above the screen. Sessions lasted approximately 10 mins.

Training trials. To encourage toddler to point systematically to the displays, they first participated in four training trials, which involved *known* words. They viewed two side-by-side dynamic scenes and were encouraged to point to one. On two trials, they were asked to point to a particular person or object in the scene (e.g., Elmo); on another two, they were asked to point to a particular

action (e.g., dancing). Toddlers who pointed to the incorrect scene were corrected gently and asked to point again.

Experiment proper. In the experiment proper, each toddler participated in six trials, each featuring a different target object and action (e.g., waving a balloon). Each trial included a Linguistic Familiarization phase, an Event Familiarization phase, and a Test phase. Trials were presented in one of two random orders, balanced across conditions. The left-right position of the test scenes was counterbalanced across trials.

Toddlers were randomly assigned to one of the two conditions, both of which featured the same visual materials; the only difference between conditions was the auditory stimuli presented in the Linguistic Familiarization phase.

Each trial proceeded as follows. First, in the Linguistic Familiarization phase (5 sec), toddlers viewed a still image taken from the first frame of the dynamic action scene (e.g., a boy waving a balloon), and heard either, “Let’s see a boy and a balloon. Let’s see *pilking*” ([+Nouns –Syntax] condition), or “A boy is gonna *pilk* a balloon. Let’s see” ([+Nouns +Syntax] condition). From this point on, toddlers are treated identically in the two conditions.

Next, in the Event Familiarization phase, toddlers viewed the dynamic action scene (12 sec). First, the scene appeared on one side of the screen, and toddlers heard “Look, *pilking!*”; after 6 sec the scene was presented on the other side of the screen, and they simply heard “Wow!”

Finally, in the Test phase, toddlers viewed two new scenes simultaneously, one on either side of the screen. One depicted the now-familiar action with a new object (e.g., waving a rake), and the other depicting the familiar object, but a new action performed on it (e.g., tapping the balloon). First toddlers heard, “Now look—they’re different!” This phase was designed to give toddlers time to inspect both scenes, as both were novel to them. The screen then went black, during which time they heard, “Do you see *pilking?*” The scenes immediately re-appeared in their original locations, with the audio, “Find *pilking!*” Toddlers’ pointing responses were recorded.

3.2 Predictions

If labeling the participants alone is sufficient to focus toddlers’ attention on the appropriate part of the scene that is being labeled by the novel verb, they should successfully map the novel verbs in both conditions. But if they require an informative syntactic context in addition to semantically informative noun phrases labeling the participants, they should only succeed in the [+Nouns +Syntax] condition.

3.3 Results

We calculated, for each toddler, the number of trials on which that toddler pointed to the Familiar Action scene (e.g., waving a rake), and divided this by

the total number of trials on which that toddler pointed. Toddlers in the [+Nouns –Syntax] condition pointed to the Familiar Action scene significantly less often ($M=49\%$) than toddlers in the [+Nouns +Syntax] condition ($M=65\%$); $t(34) = 2.4, p < .03$. Performance in the [+Nouns +Syntax] condition also differed significantly from chance (50%); $t(17) = 2.5, p < .03$. See Figure 1.




Figure 1. Experiment 1. Mean proportion of points to Familiar Action scene, expressed as a function of condition.

3.4 Discussion

When toddlers were provided with both semantically rich noun phrases labeling the verb’s arguments, and an informative syntactic context, they successfully mapped the novel verbs to the action. Especially impressive is that they extended the novel verbs to a scene preserving the action, despite the presence of a new participant object (e.g., from *waving the balloon* to *waving the rake*). In contrast, when provided with just the noun phrases labeling the arguments, but deprived of an informative syntactic context, toddlers failed to learn the novel verbs. This provides support for the hypothesis that, like adults and older children, toddlers also learn verbs more successfully when they appear in an informative syntactic context, which provides both syntactic structure and thematic-role mapping.

However, it is also possible that toddlers failed in the [+Nouns –Syntax] condition because they failed to recognize that novel word was a verb. Although the novel words appeared with verbal morphology (e.g., “Look, pilking!”), this context may not have been sufficiently informative for these young learners to determine grammatical category. We pursue this possibility in Experiment 2.

4. Experiment 2: Manipulating Semantic Richness

In Experiment 2, all toddlers were given informative syntactic contexts. We embedded the novel verbs in transitive sentences, but with pronouns labeling the arguments (e.g., “He’s pilking it”). See Mintz (2003) and Weisleder and

Waxman (in press) for evidence that pronoun contexts should help toddlers determine a new word's grammatical category.

What varied was whether they were also given semantically rich noun phrases labeling the verb's arguments. While toddlers in the [-Nouns +Syntax] condition heard just, "Let's see what happens now. He's gonna pilk it", toddlers in [+Nouns +Syntax] were given argument labels in a separate sentence: "Let's see a boy and a balloon. He's gonna pilk it".

The goals of Experiment 2 were, first, to determine whether toddlers' failure to map novel verbs in the [+Nouns -Syntax] condition in Experiment 1 was due to their failure to correctly categorize the novel word as a verb, and second, to determine the potency of syntactic frame with semantically bleached arguments as a cue to verb meaning.

4.1 Methods

Methods were identical to Experiment 1, except for the auditory stimuli presented during the Linguistic Familiarization phase. See Table 2. A different group of 30 toddlers were included for analysis (data collection in progress).

4.2 Predictions

The [+Nouns +Syntax] condition provides the same amount of information (syntactic frame, semantically rich noun phrases labeling the arguments) as the [-Nouns +Syntax] condition in Experiment 1. The only difference is how the information is packaged—in two sentences or just one. If having sufficient information is all toddlers require to succeed, then they should be able to map the novel verbs. But if they further require that information to be packaged in a particular way for them to be able to process it effectively, they may have difficulty in this condition.

In the [-Nouns +Syntax] condition, our previous results with 24-month-olds (Arunachalam & Waxman, 2009) suggest that toddlers will not be able to map the novel verbs.

4.3 Results

Performance in the two conditions did not differ significantly from each other ([-Nouns +Syntax], $N=14$, $M=46\%$; [+Nouns +Syntax], $N=16$, $M=56\%$), nor did performance in either condition differ from chance. See Figure 2.

4.4 Discussion

Toddlers failed to map the novel verbs reliably in both the [-Nouns +Syntax] and [+Nouns +Syntax] conditions. Why did toddlers fail in the [-Nouns +Syntax] condition, even though they were given as much information as in the [+Nouns +Syntax] condition in Experiment 1? The difference lay in

whether the contentful noun phrases labeling the arguments appeared in the same sentence as the novel verb and syntactic frame: In Experiment 1, the noun phrases were integrated into the sentence containing the verb, while in Experiment 2, they appeared in a separate sentence. Toddlers' differing performance in these two conditions points to a role for how linguistic information is packaged, beyond simply whether it is available.

This result also casts doubt on the interpretation that in Experiment 1, toddlers failed in the [+Nouns –Syntax] condition because they did not encode the novel word as a verb. In Experiment 2, when given both the rich semantic information and a syntactic frame that more clearly marked the novel word as a verb, toddlers still had difficulty.

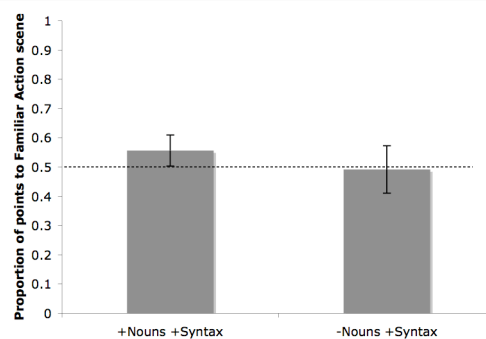


Figure 2. Experiment 2. Mean proportion of points to Familiar Action scene, expressed as a function of condition.

5 General Discussion

Taken together, these two experiments illuminate the contributions of different kinds of linguistic information to toddlers' early representations of novel verbs. We found that 27-month-old toddlers were able to map novel verbs, and further, to extend them to scenes that differed from the ones they had seen during familiarization. They were able to do so given relatively little exposure to the verbs and scenes (as compared, e.g., to Arunachalam & Waxman, 2009).

But importantly, toddlers only succeeded at mapping and extending the novel verbs when given linguistic information that was both semantically and syntactically rich. Like adults in the Human Simulation Paradigm, toddlers performed better when they were given more information ([+Nouns +Syntax]) than when given less information ([–Nouns +Syntax], [+Nouns –Syntax]).

But for toddlers, the way in which the linguistic information was packaged played an additional role. Toddlers succeeded in the [+Nouns +Syntax] condition in Experiment 1 (“A boy is gonna *pilk* a balloon. Let’s see.”), but failed in the [+Nouns +Syntax] condition in Experiment 2 (“Let’s see a boy and a balloon. He’s gonna *pilk* it.”). The presence of rich semantic and rich syntactic

information was necessary, but not sufficient, to promote successful verb mapping in this task. Only when the two kinds of information were packaged in a single sentence did toddlers succeed.

We suspect their failure in Experiment 2 is related to the processing demands of parsing the sentences and determining the antecedents of the pronouns. Processing considerations have been implicated in a recent verb learning study using similar methods, but intransitive sentences describing one-participant events (Lidz et al., 2006). Lidz and colleagues found that, contrary to Arunachalam & Waxman (2009) and the current study, toddlers performed better when verbs describing the one-participant events were presented in pronoun contexts (e.g., “It’s *pilking*”) than full noun phrase contexts (e.g., “The flower’s *pilking*”). For these one-participant events, there may be fewer possible interpretations of the scene, and so toddlers may have had less difficulty discerning the verb’s meaning. If this interpretation is correct, the extra information conveyed by the full noun phrase may have been uninformative, simply posing a processing burden.

The optimal packaging is likely to differ across languages as well. In languages which allow noun phrase arguments to be freely dropped, such that verbs often appear in isolation (e.g., Japanese, Korean), the full noun phrase sentences that we found particularly advantageous for English learners may in fact be problematic (Imai et al., 2008). Although the full noun phrases convey more information, the difficulty of processing them in languages that allow argument dropping may cancel out the benefit of informativeness. It is likely that learners perform best when an optimal balance is achieved between informativeness and ease of processing.

References

- Arunachalam, Sudha & Waxman, Sandra, R. (2009). 24-month-old infants’ use of syntactic context in noun and verb learning. Paper presented at the Society for Research in Child Development Biennial Meeting, Denver, CO.
- Fenson, Larry; Dale, Philip S.; Reznick, J. Steven; Thal, Donna J.; Bates, Elizabeth; Hartung, Jeffrey P.; Pethick, Steve; & Reilly, Judy S. (1993). User’s guide and technical manual for the MacArthur Communicative Development Inventories. San Diego, CA: Singular Press
- Gillette, Jane; Gleitman, Henry; Gleitman, Lila R.; & Lederer, Anne. (1999). Human simulations of vocabulary learning. *Cognition*, *73*, 135–176.
- Gleitman, Lila R. (1990). The structural sources of verb meanings. *Language Acquisition*, *1*, 1-55.
- Imai, Mutsumi; Li, Lianjing; Haryu, Etsuko; Okada, Hiroyuki; Hirsh-Pasek, Kathy; Golinkoff, Roberta; & Shigematsu, Jun. (2008). Novel noun and verb learning in Chinese-, English-, and Japanese-speaking children. *Child Development*, *79*, 979-1000.
- Lidz, Jeffrey; Bunger, Ann; Leddon, Erin; & Waxman, Sandra R. (2006). When one cue is better than two: Syntactic v. lexical information in infant verb learning. Talk presented at the CUNY Workshop on On-Line Methods in Children’s Language Processing, New York.

- Mintz, Toben. (2003). Frequent frames as a cue for grammatical categories in child directed speech. *Cognition*, 90, 91-117.
- Naigles, Letitia. (1990). Children use syntax to learn verb meanings. *Journal of Child Language*, 17, 357-374.
- Snedeker, Jesse; & Gleitman, Lila R. (2004). Why it is hard to label our concepts. In *Weaving a Lexicon*. Cambridge: MIT Press.
- Waxman, Sandra R.; Lidz, Jeffrey; Braun, Irena E.; Lavin, Tracy. (2009) Twenty-four-month-old infants' interpretations of novel verbs and nouns in dynamic scenes. *Cognitive Psychology*, 59, 67-95.
- Weisleder, Adriana; & Waxman, Sandra R. (in press). 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*.
- Yuan, Sylvia; & Fisher, Cynthia. (2009). Really? She blicked the baby? Two-year-olds learn combinatorial facts about verbs by listening. *Psychological Science*, 20, 619-626.