Combining Syntactic Frames and Semantic Roles to Acquire Verbs

Rose M. Scott and Cynthia Fisher
University of Illinois

For any given utterance of a verb, the referential scene offers a wide array of potential interpretations. The syntactic bootstrapping hypothesis (Landau & Gleitman, 1985) maintains that children could constrain these interpretations by exploiting systematic links between syntactic structure and verb meaning. A number of studies have provided support for this hypothesis. Children interpret a novel verb in a transitive sentence as referring to a relationship between two participants (Fisher, 2002; Naigles, 1990; Yuan & Fisher, 2006). Children also draw upon language-specific syntactic knowledge, such as word order in English, to map a novel verb to an appropriate event (Gertner, Fisher, & Eisengart, 2006). These findings suggest that the syntactic structure of an utterance can act as a zoom lens, limiting children's interpretation of a verb to a particular aspect of the referential scene.

However, the constraints from a single frame provide only highly abstract information about the meaning of a verb. To illustrate, consider encountering a novel transitive verb, as in "John blicked the doctor." Occurrence in the transitive frame indicates that *blick* describes a relationship between two participants. A wide range of meanings are consistent with this constraint: *blicked* could mean anything from *killed* or *murdered* to *loved* or *visited*. Since many diverse meanings can map onto this single frame, the meaning of *blick* remains highly ambiguous.

How do children overcome this ambiguity? Observation of a verb's use across multiple scenes is one likely source of further constraint. However, Landau and Gleitman (1985) proposed that children could also use the set of frames a verb occurred in as an additional source of constraint on its meaning. For instance, the verb *explain* occurs in both a dative frame (1a) and in a sentential complement frame (1b). Each of these frames provides different constraints on the meaning of *explain*. Occurrence with direct and indirect objects (e.g., in a dative construction) suggests that the verb describes transfer of possession or motion towards a goal. Occurrence with sentential complements indicates that the verb denotes a predicate involving propositional content. Together, these constraints considerably narrow the potential meaning of *explain*: it involves the transfer of propositional content.

(1) a. Mary explained the problem to me.
b. John explained that she left.

Work by Naigles (1996, 1998) suggests that children can use the set of frames in which a verb occurs to constrain its interpretation. Specifically, she examined the causal alternation (2) and the unspecified-object alternation (3). Verbs that occur in the causal alternation (2) describe events with the complex internal structure of an action with a result sub-event (e.g., externally-caused motion or change-of-state). In contrast, verbs that participate in the unspecified-object alternation (3) describe activities (Levin, 1993) and do not entail any particular effect on the target object.

(2) a. Anne broke the lamp.
b. The lamp broke.

(3) a. Anne dusted the lamp.
b. Anne dusted.

Naigles (1996, 1998) presented children with these alternations in a preferential-looking task. Children simultaneously viewed a caused-motion event (e.g., a duck pushing on a frog's shoulders, causing the frog to bend) and a contact-activity event (a duck patting a frog on the head). The complex internal structure of the caused-motion event makes it an appropriate referent for the causal alternation whereas the activity nature of the contact-activity event makes it an appropriate referent for the unspecified-object alternation. While viewing these events, children heard a novel verb (e.g., *sebbing*) presented in either the causal or the unspecified-object alternation. Results from control conditions indicated a baseline preference for the caused-motion event, making it difficult to draw conclusions about the effect of the causal alternation. However, an encouraging difference emerged in the unspecified-object condition: children looked equally at the two events. Thus, hearing the verb used in the unspecified-object alternation influenced their interpretation of the novel verb, causing them to look relatively more at the contact-activity event than they would have otherwise. These results suggest that children can use these two sets of frames to interpret verbs they encounter.

Although Naigles' (1996) results suggest that children can use multiple frames to interpret verbs, the actual mechanism behind multiple-frames bootstrapping remains unclear. Pinker (1994) put forth one possibility: the multiple frames bootstrapping process might operate via the use of language-specific verb classes. Languages typically have classes of verbs that demonstrate hyper-similarity in both their syntactic and semantic behavior. The grouping of verbs into classes is highly arbitrary and language-specific, meaning they must be learned. Pinker proposed that when children encounter a verb in a set of frames for which they have established a class, they extend to it the shared meaning of the other verbs in the class (see Pinker, Lebeaux, & Frost, 1987 for evidence of this in older children). Given that verb classes are language-specific, the use of these classes to draw inferences about novel verbs would require considerable prior learning about the lexicalization patterns of a particular

language. If multiple-frames bootstrapping operates solely via verb classes, the necessity of this prior learning would effectively rule it out as an early verb learning mechanism.

An alternative proposal by Fisher, Hall, Rakowitz, and Gleitman (1994) suggested that multiple-frames bootstrapping could occur via an iterative application of the single-frame bootstrapping process. When children encountered a verb, they would make an inference about that verb's meaning based on the constraints provided by the syntactic frame in which it occurred. Upon encountering that verb in another frame, they would make a new semantic inference consistent with that new frame. If children then took into account the constraints provided by both hypothesized meanings, they could arrive at a more precise verb meaning. This process could be repeated for additional frames, with each frame resulting in a further refinement of the verb's meaning. Note that this process requires less language-specific learning than does the use of verb classes. Having acquired just a few nouns, children could use general mappings between syntactic and conceptual structures to make inferences about verb meaning. Given prior evidence that these general constraints are available to very young children, this makes iterative-processing of multiple frames a plausible mechanism for early verb acquisition.

However, the effectiveness of the iterative-processing mechanism does depend on how children represent utterances. If children merely represented the surface structure of the sentence, they would form representations such as (4). Based on these representations, children would be able to infer that *break* could describe either a two-participant relation (4a) or a one-participant predicate (4b). This leaves the meaning of *break* ambiguous and does not distinguish it from non-causal verbs such as *dust*, which also describe both two-participant and one-participant predicates. If children only represented surface structure, iteratively processing sets of frames would not help them to refine verb meanings.

(4) a. Anne broke the lamp. → NP verb NP.
b. The lamp broke. → NP verb.

Recent evidence suggests that children can track more than just the surface structure of a sentence. Two-year-olds can use verbs' semantic selection restrictions to anticipate potential direct objects (Chang & Fernald, 2003) and to acquire novel terms (Goodman, McDonough, & Brown, 1998). These findings indicate that young children can track specific semantic information about the arguments of verbs. Given this capacity, children could annotate surface structures with rough semantic roles, creating representations such as those in (5). Using these representations, children could infer that *break* can describe a causal action (5a) or just the effect on the object (5b). The fact that this effect is worthy of its own sentence suggests that it comprises a separate event. Combined with the transitive frame, this implies that the referent event is one with the complex internal structure of causal action + result subevent.

(5) a. Anne broke the lamp. → NP verb NP
<actor> <undergoer>
b. The lamp broke. → NP verb.
<undergoer>

Frame representations annotated with rough semantic information appear to provide the correct type of information for 1) making useful inferences based on a single frame and 2) combining inferences across multiple frames. This paper therefore proposes an elaborated version of Fisher et al.'s (1994) iterative-processing model of multiple-frames bootstrapping. This model assumes that children semantically annotate verb frames as they encounter them. They then use semantically-annotated structures to draw inferences about verb meanings. Children form new inferences for each frame they encounter and use these new inferences to refine their interpretation of a verb, making multiple-frames bootstrapping an iterative process.

2. Sources of role-relevant information

The iterative-processing mechanism operates on semantically-annotated frames and thus depends on children's ability to identify the likely semantic roles of participants in sentences. For instance, in order to make useful inferences about the meaning of verbs in the causal and unspecified-object alternations, children would need to notice that the two verb types assigned different semantic roles to the subject position of their respective intransitive frames. Merlo and Stevenson (2001) argued that this underlying difference could be detected using surface features of the input. They found that within the Wall Street Journal corpus, causal and unspecified-object verbs could be distinguished with 69% accuracy using primarily three features: transitivity, subject animacy, and lexical overlap between subject and object position.

Recently, Scott and Fisher (2006) extended these findings to child-directed speech. They examined parental utterances containing the most frequently occurring causal and unspecified-object verbs in a large corpus of child-directed speech obtained from the CHILDES database (MacWhinney, 2000). Target verbs were coded for transitivity, overall subject animacy, intransitive-subject animacy, and object animacy. Classification analyses using all variables resulted in 90% classification accuracy, with the strongest predictor of verb class being intransitive-subject animacy. These findings indicate that the surface features of child-directed speech reliably reflect the underlying argument structures of causal and unspecified-object verbs. This suggests that children have access to at least one source of information that would permit them to form the semantically-annotated structures required for iterative-frame processing.

Given that children have access to a source of role-relevant information, the iterative-processing model predicts that children should use this information in conjunction with syntactic structure to make inferences about novel verbs. In order to test this prediction, Experiment 1 drew on a second potential source of

role-relevant information: conceptual representations of referential scenes. Consider the events in Figure 1. In the caused-motion event (left), the girl is pushing down on the boy's shoulders, causing him to bend at the knees in a squatting motion. Children should represent this event in terms of two subcomponents: the girl bending the boy and the boy's ensuing bending action. In the contact-activity event (right), the girl is dusting the boy's back with a feather duster. Here, children should represent the girl dusting the boy as well as just the girl's dusting action. While viewing these events, children were presented with a novel verb in either the causal alternation (e.g., "The girl is pimming the boy. The boy is pimming.") or the unspecified-object alternation (e.g., "The girl is pimming the boy. The girl is pimming."). If children map their conceptual representations onto the stimulus sentences to make inferences about verbs, they should identify the caused-motion event as a suitable referent for a novel verb used in the causal alternation. Similarly, they should map a novel verb in the unspecified-object alternation to the contact-activity event.

Figure 1. Caused-motion (left) event and contact-activity event (right).

The iterative-processing model also predicts that, in addition to making inferences using sets of frames, children should be able to make useful inferences about verb meaning based on a single semantically-annotated frame. To test this prediction, additional groups of children were presented with either the intransitive frame of the causal alternation ("The boy is pimming") or the intransitive frame of the unspecified-object alternation ("The girl is pimming"). If children reason about individual frames in conjunction with the roles provided by their conceptual representations, then the sentence "The boy is pimming" should direct their attention to the caused-motion event, as the result sub-event provides a subject-worthy role for the boy. In contrast, the contact-activity event does not provide the boy with a subject-worthy role. Since both events supply the girl with an actor role, the sentence "The girl is pimming" could apply to both events.

3. Experiment 1

3.1. Methods

3.1.1. Participants

Forty 28-month-old children (20 male, 20 female; mean age 28.0 months, range 27.0-29.9 months) participated in a preferential looking task. All children were native speakers of English. Eight children were randomly assigned to each

of the five soundtrack conditions as shown in (6) below. Three additional children were tested but excluded from analyses due to side bias (more than 80% of practice trials spent looking to one side; n=1) or looking times greater than 2.5 SD away from the condition mean (n=2). Children's productive vocabulary was measured using the short form of the MacArthur CDI, Level II (Fenson, Pethick, Renda, Cox, Dale, & Reznick, 2000). Scores ranged from 23 to 96 with a mean of 72. Children were randomly assigned to one of five conditions: causal, undergoer-subject, neutral, actor-subject, and unspecified-object. Each condition had an equal number of boys and girls. Mean age and vocabulary did not differ between conditions.

3.1.2. Apparatus

Children sat on a parent's lap viewing two 20-inch television monitors located 30 inches away. The screens were separated by a 12-inch gap and placed at eye level. Soundtracks were presented from a central speaker. A hidden camera recorded the children's eye movements throughout the experiment. Parents wore opaque sunglasses, preventing them from biasing their children's responses to the videos.

3.1.3. Procedure

Children viewed a synchronized pair of videos involving a boy and a girl. The videos were accompanied by a soundtrack recorded by a native English speaker. The video sequence consisted of three phases: familiarization, practice, and test. In the familiarization phase, children first viewed an event in which one screen displayed the girl waving while the other screen remained blank (4s) and the girl was labeled twice ("There's a girl!"). This was followed by a 2-second interval during which both screens were blank. Children then viewed an event in which the boy waved on one screen (4s) while being labeled twice ("There's a boy!") and other screen remained blank. These events were then followed by two 4-second trials in which the girl appeared on one screen and the boy on the other. In one trial, children were asked to "Find the boy!" and in the next they were asked to "Find the girl!"

The familiarization phase was followed by a brief practice phase involving two familiar intransitive verbs. The practice phase began with a 4-second interval in which both screens were blank. During this interval, the first practice action was labeled in the future tense ("The boy's gonna jump!). Children then viewed an 8-second trial in which they saw the boy jumping on one screen and pretending to sleep on the other; the soundtrack labeled the boy's jumping action three times ("The boy is jumping!"). This was followed by another 4-second interval of blank screen, during which the jumping action was labeled once in the past tense ("The boy jumped!") and children were instructed to "Find jumping." The children viewed the jumping and sleeping events again (8s) and were asked to "Find jumping." The second set of practice trials was identical to

the first set except that children now viewed the girl clapping on one screen and eating on the other and were asked to "Find clapping!"

The familiarization and practice phases were intended to familiarize the children with the actors. These trials also served to familiarize the children with the fact that two events would be presented on each trial and the soundtrack would ask them to look at one of the two events.

The pair of videos shown in the test phase was the caused-motion and contact-activity events depicted in Figure 1. During the test phase, the audio varied according to condition, such that children heard one of the following 5 sets of sentences shown in (6):

(6) Causal: "The boy is pimming. The girl is pimming the boy."
Unspecified-object: "The girl is pimming. The girl is pimming the boy."
Undergoer-subject: "The boy is pimming. The boy is pimming."
Actor-subject: "The girl is pimming. The girl is pimming."
Neutral: "Look! Isn't that fun? See?"

Table 1. Test phase, shown with audio from the causal condition.

Duration	Left screen	Right screen	Audio
7 sec.	<Blank screen>	<Blank screen>	"The girl is gonna pim the boy. The boy's gonna pim!"
8 sec.	Girl bending boy.	Girl dusting boy.	"The boy is pimming. The girl is pimming the boy. See?"
7 sec.	<Blank screen>	<Blank screen>	"The girl pimmed the boy. The boy pimmed! Find pimming."
8 sec.	Girl bending boy.	Girl dusting boy.	"The boy is pimming. The girl is pimming the boy. Find pimming."

Table 1 depicts a sample test phase using audio from the causal condition. For all children, the test phase was separated from the practice phase by a 7-second blank-screen interval. During this interval, children in the causal, unspecified-object, undergoer-subject, and actor-subject conditions heard the novel verb *pimming* used twice in the future tense, in the sentence structures appropriate to their condition. Children in the neutral condition heard neutral audio (e.g., "What's gonna happen?"). This was followed by an 8-second test

trial in which children viewed the test events while hearing the pair of sentences for their condition, this time in the present progressive. The test trial was followed by a second blank-screen interval during which the sentences were repeated in the past tense and children were asked to "Find pimming" (with the exception of those in the neutral condition, who heard "Did you like that?"). This was followed by a second 8-second test trial in which the pair of sentences (or the neutral audio) were repeated. Presentation side of the caused-motion event was counterbalanced within each condition and each sex.

3.1.4 Coding

Children's visual fixations were coded frame by frame from silent video. The visual fixations of 25% of the children were assessed by a second coder. The two coders agreed on the direction of the child's gaze in 98% of the frames. For each trial, the time spent looking at the caused-motion event was divided by the total time spent looking to the two screens. The proportion of looking-time to the caused-motion event was then averaged across the two test trials. Any trial in which the child looked away for more than 50% of the trial was eliminated (n=1).

3.2 Results

Table 2 shows the mean proportion of looking time to the caused-motion event, separately by condition. Preliminary analyses revealed no effects of sex, or whether vocabulary or practice performance was above or below the median. These factors were therefore not examined further.

Table 2. Mean (SD) proportion of looking-time to the caused-motion event

Condition	Looking-time proportion
Causal alternation	.55 (.10)
Unspecified-object alternation	.35 (.16)
Undergoer-subject intransitive	.54 (.11)
Actor-subject intransitive	.41 (.11)
Neutral	.39 (.08)

Two separate series of analyses were conducted. The first set of analyses asked whether children assigned different interpretations to the novel verb presented in the causal and unspecified-object alternations. As Table 2 shows, children who heard the novel verb used in the causal alternation looked longer at the caused-motion event than did those who heard either neutral audio or the

novel verb used in the unspecified-object alternation. This pattern of results was supported by an ANOVA comparing the proportion of looking-time to the caused-motion event in the causal, neutral, and unspecified-object conditions, which revealed a significant effect of condition ($F(2,21) = 6.427$, $p < .01$). Planned, one-tailed comparisons were conducted: causal vs. unspecified-object, causal vs. neutral, and neutral vs. unspecified-object. The comparison of the causal and unspecified-object conditions revealed that children in the causal condition looked significantly longer at the caused-motion event than did those in the unspecified-object condition ($t(21) = 3.393$, $p < .005$). Children in the causal condition also looked significantly longer at the caused-motion event than did those in the neutral condition ($t(21) = 2.699$, $p < .01$). Looking times for the unspecified-object and neutral conditions did not differ ($t(21) = .695$, NS).

A second set of analyses asked whether children could interpret the novel verb using only a single sentence frame in conjunction with the role information from the referential scene. Table 2 shows that children in the undergoer-subject condition, like those in the causal condition, looked longer at the caused-motion event than did those in the actor-subject or neutral condition. An ANOVA examining the proportion of looking-time to the caused-motion event in the undergoer-subject, neutral, and actor-subject conditions again found a significant effect of condition ($F(2,21) = 4.483$, $p < .05$). Planned comparisons revealed that children in the undergoer-subject condition looked significantly longer at the caused-motion event than did those in the actor-subject condition ($t(21) = 2.363$, $p < .05$) or the neutral condition ($t(21) = 2.774$, $p < .01$). Children in the actor-subject condition did not differ from those in the neutral condition ($t(21) = .411$, NS).

4. Discussion

The results of this experiment extend the findings of Naigles (1996, 1998), demonstrating that 28-month-old children can use their conceptual representations of events to make inferences about verbs presented in multiple frames. Specifically, children who heard a novel verb used in the causal alternation looked longer at the caused-motion event than did those who heard it used in the unspecified-object alternation. In addition, children who heard the verb presented in the causal alternation looked longer at the caused-motion event than did those in the neutral condition. This clarifies Naigles' results by establishing that children can make inferences about a novel verb based on its occurrence in the causal alternation. Children who heard the unspecified-object alternation failed to exhibit a significant preference for the contact-activity event. This may have resulted from the baseline preference for the contact-activity event exhibited by the children in the neutral condition[1]. Given a set of

1. Alternatively, children may have construed the verb as referring to a contact-activity component of the caused-motion event (e.g., "The girl is pushing the boy. The girl is pushing.").

events with a different baseline, children may display an effect of the unspecified-object alternation on verb interpretation (e.g., Naigles, 1996). Together with Naigles' results, these findings show that children can make inferences about novel verbs that occur in either the causal or unspecified-object alternation.

The results in the intransitive conditions of this experiment showed that children can make inferences about a novel verb using individual frames in conjunction with the semantic-role information present in the referential scene. The iterative-processing model predicted that if children could draw upon the roles present in their conceptual representations of the event, then in the referential context provided they should be able to make inferences about the verb meaning using only a single frame of the causal alternation. As predicted, when presented with the undergoer-subject sentence, children looked longer at the caused-motion event than at the contact-activity event. Although both events contained a boy, children did not interpret the verb as referring to the boy in the contact-activity event. This means that they considered the boy's actions in the contact-activity event (i.e. smiling, head turning, passive recipient of action) to be poorer referents for the verb than his actions in the caused-motion event were (i.e. bending). In order for children to determine which of these actions constituted a likely verb meaning, they must have represented the event in terms of roles played by each participant (e.g., (bend(boy))) and then ranked those roles according to some criteria (e.g., salience of motion, etc.). Children selected the boy's bending action as the referent of the verb because it outranked smiling and passive recipient of action, making it the most subject-worthy role. Having mapped the verb onto the boy's bending action, children could infer that the subject of the verb was an undergoer, as this was the role played by the boy in the caused-motion event. If children were to annotate the surface structure of the sentence with this role information, they would be able to use it to constrain verb interpretations in the future.

The results of the undergoer-subject condition support the iterative-processing model's prediction that children can make useful inferences about the meaning of a verb based on each sentence frame in which they encounter it. This finding supports an alternative to the class-based multiple-frame bootstrapping mechanism presented by Pinker (1994). On Pinker's view, children must encounter the entire alternation and then assimilate the verb to an existing learned class before they can draw useful conclusions about the verb's meaning based on a set of frames. It is possible that 28-month-old children have acquired the causal and unspecified-object verb classes for English and that the performance in the causal and unspecified-object conditions reflects the use of these classes. However, the fact that children in the undergoer-subject condition preferentially mapped the verb to the caused-motion event based on just a single frame suggests that the use of learned verb classes is not *required* for making helpful inferences about the verbs that occur in these alternations. By considering the syntactic structure in conjunction with the roles of the participants in the sentence, children were able to make inferences based on a

single frame. This capacity would permit multiple-frames bootstrapping to aid children early in the language acquisition process, prior to the acquisition of language-specific verb classes.

The results of the undergoer-subject condition give rise to a potential complication: both the causal alternation and the undergoer-subject sentence caused children to map the verb to the caused-motion event. One could interpret this result as suggesting that sets of frames do not contribute any more information than individual frames do, as the intransitive alone seems to provide sufficient information for arriving at the correct interpretation of the verb. However, while children in both conditions did map the verb onto the same event, this does not mean that they arrived at the same interpretation of the verb. Bunger and Lidz (2006) found that when children heard a verb used in the undergoer-subject frame, they interpreted it as referring to the result sub-component of a caused-motion event and did not include the causal component in their interpretation of the verb. In this case, children who only heard the undergoer-subject sentence would require experience with the transitive frame to determine that the verb also encodes the cause of the undergoer's motion. Thus, while the undergoer-subject frame supports useful inferences about the verb's meaning, both frames are required to arrive at the full meaning of the verb.

5. Conclusion

This paper proposed an iterative-processing model of multiple frames bootstrapping. The model states that children could make inferences based on each frame they encounter, iteratively refining verb meanings based on new verb-frame pairings. It was argued that in order for this process to work, children would need to represent more than merely the surface structure of the sentence. Specifically, it was proposed that children annotate sentences with semantic-role information. In conjunction with Scott and Fisher (2006), the findings reported here demonstrate that children have access to two sources of information about semantic roles: the surface properties of the input (e.g., subject animacy) and their conceptual representations of the referential scene. Furthermore, the results of the current experiment show that they can use the latter to reason about verbs presented in the causal and unspecified-object alternations. This experiment demonstrated that children can interpret novel verbs using a single sentence frame in conjunction with semantic-role information, here provided by the referential scene. This provides some initial support for the basic mechanism proposed by the iterative-processing model. Future experiments will need to examine directly whether children refine verb interpretations iteratively over time as proposed by the model.

References

Bunger, A., & Lidz, J. (2006). Constrained flexibility in the acquisition of causative verbs. *Proceedings of the Annual Boston University Conference on Language Development, 30,* 60–71.

Chang, E., & Fernald, A. (2003). Use of semantic knowledge in speech processing by 26-month-olds. *Monographs of the Society for Research in Child Development, 68.*

Fenson, L., Pethick, S., Renda, C., Cox, J.L., Dale, P.S., & Reznick, J.S. (2000). Short form versions of the MacArthur Communicative Development Inventories. *Applied Psycholinguistics, 21,* 95-115.

Fisher, C. (2002). Structural limits on verb mapping: The role of abstract structure in 2.5-year-olds' interpretations of novel verbs. *Developmental Science, 5,* 56-65.

Fisher, C., Hall, G., Rakowitz, S., & Gleitman, L.R. (1994). When it is better to receive than to give: Syntactic and conceptual constraints on vocabulary growth. *Lingua, 92,* 333-375.

Gertner, Y., Fisher, C., & Eisengart, J. (2006). Learning words and rules: Abstract knowledge of word order in early sentence comprehension. *Psychological Science, 17,* 684-691.

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,* 1330-1344.

Landau, B., & Gleitman, L. R. (1985). *Language and experience: Evidence from the blind child.* Cambridge, MA: Harvard University Press.

Levin, B. (1993). *English verb classes and alternations: A preliminary investigation.* Chicago, IL: University of Chicago Press.

MacWhinney, B. (2000). *The CHILDES project: Tools for analyzing talk.* Third Edition. Mahwah, NJ: Lawrence Erlbaum Associates.

Merlo, P., & Stevenson, S. (2001). Automatic verb classification based on statistical distributions of argument structure. *Computational Linguistics, 27,* 373-408.

Naigles, L. (1990). Children use syntax to learn verb meanings. *Journal of Child Language, 17,* 357-374.

Naigles, L. R. (1996). The use of multiple frames in verb learning via syntactic bootstrapping. *Cognition, 58,* 221-251.

Naigles, L. R. (1998). Developmental changes in the use of structure in verb learning: evidence from preferential looking. *Advances in Infancy Research, 12,* 298-317.

Pinker, S. (1994). How could a child use verb syntax to learn verb semantics? *Lingua, 92,* 377-410.

Pinker, S., Lebeaux, D.S., & Frost, L.A. (1987). Productivity and constrains in the acquisition of the passive. *Cognition, 26,* 195-267.

Scott, R. M., & Fisher, C. (2006). Automatic classification of transitivity alternations in child-directed speech. In *Proceedings of the 28th Annual Meeting of the Cognitive Science Society*, 2129-2134.

Yuan, S., & Fisher, C. (2006). "Really, he blicked the cat?": 2-year-olds learn distributional facts about verbs in the absence of a referential context. *Proceedings of the Annual Boston University Conference on Language Development, 30,* vol. 2, 689-700.