Automatic Classification of Transitivity Alternations in Child-Directed Speech

Rose M. Scott (rmscott2@uiuc.edu)
Department of Psychology, 603 E. Daniel Street
Champaign, IL 61820 USA

Cynthia Fisher (clfishe@uiuc.edu)
Department of Psychology, 603 E. Daniel Street
Champaign, IL 61820 USA

Abstract

Children can use the characteristic entailment patterns of verb classes to learn new verbs (Naigles, 1996), but how do they acquire verb classes? One possibility is that surface features that reflect argument structure can separate verb classes (Merlo & Stevenson, 2001). We tested this hypothesis by running clustering analyses on surface features of 29 verbs extracted from a large corpus of child-directed speech. Results suggest that animacy plays an important role in verb classification.

Introduction

The syntactic bootstrapping hypothesis (Gleitman, 1990; Landau & Gleitman, 1985) holds that children use syntax to guide verb learning. For a given utterance of a verb, the referential scene offers a plethora of potential meanings. Given principled relationships between clause syntax and verb meaning (e.g., Grimshaw, 1990; Jackendoff, 1990; Levin & Rappaport-Hovav, 2005), the syntactic structure, or frame, in which a verb occurs can act as a kind of linguistic "zoom lens" that constrains those interpretations. This hypothesis has been supported by many studies showing that children as young as 21 months old assign different meanings to verbs presented in different sentence structures.

For example, if children hear a new verb in a transitive sentence, they interpret the verb as referring to the action of one participant on another; if the verb is intransitive, they interpret it as referring to an event requiring only one participant (e.g., Fisher, 2002; Naigles, 1990; Yuan & Fisher, in press). Furthermore, each advance in learning the syntactic and morphological features of the native language yields new constraints on verb learning. For example, English-learning infants as young as 21 months old use word order in a transitive sentence to guide interpretation of a new verb (Gertner, Fisher, & Eisengart, in press).

A single sentence frame, however, is a limited guide to verb meaning.

First, it provides only highly abstract semantic information. The transitive frame occurs with such disparate meanings as *He became a doctor* and *He shot a doctor*. The verbs that appear in this frame share, not the specifics of the events they describe, but similar formal structure: becoming and shooting both require two core arguments (though only one person in the case of *become*). A sentence frame yields information mainly about the number and type of arguments associated with the verb–what Grimshaw (1994) has termed its semantic structure–rather than its semantic content. Observation of events must provide the semantic content (see Fisher, 2000, for a review). Experimental results suggest that 2-year-olds are appropriately open-minded about the possible meanings of transitive verbs: when children are presented with a novel verb in a transitive sentence, they can interpret it as referring to a caused-motion event (e.g., pushing) or to a simple contact event (e.g., patting; Naigles & Kako, 1993).

Second, a single frame yields information only about the semantic structure of the verb when used in that frame. Most verbs occur in more than one frame, each of which results in a different semantic structure, as in (1).

(1) a. She explained that he left.
    b. She explained the problem to me.

Gleitman and colleagues have proposed that children can gain further constraint on verb learning by appealing to the set of subcategorization frames in which a verb appears (Gleitman et al., 2005). For example, *explain* occurs with sentence complements (1a); this frame implies that one of the verb's arguments has propositional content. *Explain* also occurs with both direct and indirect objects (1b), a frame consistent with transfer, or motion toward a goal. The combination of these frames considerably narrows the possible meanings of *explain*: it is a verb describing transfer of propositional content.

But even a set of subcategorization frames can be ambiguous. Take for instance the verbs *break* (2) and *dust* (3):

(2) a. Anna broke the lamp.
    b. The lamp broke.
(3) a. Anna dusted the lamp.
    b. Anna dusted.

Both can be transitive and intransitive. Given only their occurrence in these two frames, a learner would have to treat both verbs as members of a single syntactically defined class and use world observation to infer the semantic content of each verb in the class. The learner who does this, however, would be missing a very useful distinction between verbs like *break* and verbs like *dust*, which can be seen if the entailment relations across the transitive and intransitive sentences in each pair are taken into account.

The transitive sentence (2a) describes an event with two parts: the application of force to the lamp and the ensuing breaking sub-event. This sentence necessarily entails the sub-event (that the lamp broke), expressed in the intransitive sentence (2b). In contrast, sentence (3a) does not describe a

complex event and entails nothing much about the lamp. Instead, it entails that the agent performed the dusting activity, expressed in the intransitive sentence in (3b).

The argument-structure alternation shown in (2) is known as the causal alternation; verbs participating in this alternation have the complex internal structure of a causal event with a result sub-event (e.g., Levin, 1993; Pinker, 1989). The alternation in (3) is the unspecified-object alternation; verbs participating in this alternation include many activity verbs (e.g., Levin, 1993).

Experimental evidence suggests that by 28 months of age, children can use the characteristic entailment patterns of these two verb classes to help identify the referent of a novel verb (Naigles, 1996; Scott & Fisher, in progress). Given that the two syntactic frames involved in these two alternations are the same, how do children learn that these are two classes of verbs?

Merlo and Stevenson (2001, henceforth M&S) proposed that surface properties of the input could be used to classify verbs as occurring in either the causal or unspecified-object alternation[1]. Using text from the *Wall Street Journal* (WSJ), they found that the two classes of verbs could be identified with reasonable accuracy (>69%) based primarily on three features: frequency of occurrence in the transitive frame (transitivity), subject noun phrase animacy, and lexical overlap between the subject and object positions (causativity).

Note in examples (2-3) that two of these features in particular (animacy & causativity) should directly reflect the underlying argument structure of these verbs: Unspecified-object verbs (3) assign the same thematic role (agent) to their subjects regardless of transitivity. Causal verbs (2), assign the same role (theme) to the object of the transitive and to the subject of the intransitive. The causativity measure was designed to estimate this thematic-role overlap. Similarly, the agent subjects of unspecified-object verbs should tend to be animate; causal verbs, in contrast, with theme subjects in intransitive uses, should have more inanimate subjects. M&S's results suggest that these features could play an important role in identifying verbs that occur in the unspecified-object and causal alternations, thus automatically dividing verbs into usefully narrow semantic sets.

The WSJ corpus on which these results were based, however, differs greatly from casual speech to children. Comparisons of newspaper text and adult-directed conversational speech yield different estimates of the same verbs' subcategorization probabilities, in part because verbs tend to be used in different senses in different discourse styles (e.g., Roland & Jurafsky, 1998). Speech to children also differs in many ways from conversation among adults: Child-directed speech is characterized by short sentences, repetitiveness, and simplified vocabulary, for example (e.g., Bard & Anderson, 1994; Newport, Gleitman & Gleitman, 1977).

Could children use transitivity, animacy, and causativity to classify verbs? To find out, we examined the distribution of these and related features in a corpus of child-directed speech (CDS), and used an unsupervised clustering algorithm to determine whether these distributions differentiated causal and unspecified-object verbs in CDS.

Methods

Materials

First, we searched the CHILDES database of transcribed conversations with children (MacWhinney, 2000) for part-of-speech tagged corpora containing parental utterances to target children less than 30 months of age. The following 10 corpora were selected based on these criteria: Bloom 1970, Brown, Clark, Demetras Working, Higginson, Kuczaj, New England, Post, Suppes, and Warren-Leubecker. These corpora contained 112,000 parental utterances. Next, we identified the subset of these utterances that contained verbs, using CLAN tools for searching CHILDES transcripts and relying on the existing part-of-speech tagging as of 2/5/05. Finally, we selected the verbs that occurred more than 30 times in total and that appeared in more than 5 different corpora. Among the resulting set of frequent verbs, we identified 15 unspecified-object verbs and 14 causal verbs; these are listed in Table 1 (classification based primarily on Levin, 1993). The 12,521 utterances containing these verbs were coded using the procedures described below.

Table 1: Verbs used in the clustering experiments

Verb class	Selected verbs
Understood-Object	bite, draw, drink, eat, hit, play, pull, push, read, see, throw, tickle, try, wash, write
Causal	bounce, break, change, close, fold, move, open, pop, roll, shut, slide, spill, tear, turn

Coding

The coding was first done using the CLAN program and the existing part-of-speech tagging. General search heuristics were used to identify potential subjects and objects. These heuristics also limit our assumptions about children's ability to track arguments across long, complex utterances. To ensure accuracy, all coded utterances were later hand checked by the first author.

Transitivity An utterance was coded as transitive if the verb was followed by a noun, pronoun, determiner, or any of a set of quantifiers (*some, any, all, much, more*). An utterance was coded as intransitive if the verb was followed by a punctuation mark, a conjunction, a preposition[2], a locative phrase, another verb, or a filler (e.g., *uh-oh, huh*). This heuristic coding assumed that children did not yet understand the function of "wh" words, thus coding sentences such as "What did you break?" as intransitive utterances. In a small proportion of the utterances the verb was followed by a part of speech other than those mentioned above (e.g., adjectives/adverbs). These were not coded, as pilot analyses

---

[1] M&S's study also included a third class, induced-action verbs that take a form of the causal alternation in which both arguments are animate (e.g., *He jumped the horse over the fence*).

[2] Utterances containing phrasal verbs (i.e "tore up the paper") were hand corrected to be transitive rather than intransitive.

showed they did not reliably fall into either the transitive or intransitive category. The final data set, after transitivity coding, consisted of 11,780 utterances.

Animacy Each utterance was coded for subject animacy. In addition to the overall subject animacy measure used by M&S, we calculated separate animacy scores for transitive and intransitive subjects. We also coded direct-object animacy. This addition was intended to approximate the causativity measure used by M&S. We reasoned that if causal alternation verbs exhibit both lower overall subject animacy and greater subject/object lexical overlap than unspecified-object verbs, as M&S found, then causal verbs should have fewer animate direct objects than unspecified-object verbs.

For these initial analyses, we used pronouns as an automatically extractable approximation of animacy (see M&S for a similar procedure). Other analyses make clear that pronoun arguments are very common in child-directed speech; moreover, the distribution in sentences of particular pronouns can be used to sort verbs into different semantic classes (Laakso & Smith, 2004), and also plays an important role in defining distributional 'frequent frames' that can be used to distinguish verbs from nouns (Mintz, 2003). Pronoun arguments that signal an animacy contrast between subject and object also make it easier for children to parse transitive sentences, and to generalize the transitive structure across new verbs in a training study (Childers & Tomasello, 2001). Given all this, it seems reasonable to assume that children know the meanings of many pronouns from an early age.

An utterance was coded as having an animate subject if the verb was preceded by (permitting one intervening auxiliary) *he, she, we, I, you, let's* or *who*. Verb-initial utterances (i.e. imperatives) were also coded as having animate subjects. Inanimate subjects were *it, that, this, that one, this one*, or *what*. An utterance was coded as having an animate object if the verb was followed by *him, her, us, me, you*, or *who(m)*. If the verb was followed by *it, that, this, that one, this one*, or *what*, this indicated an inanimate object. These heuristics captured 75% of the subjects and 43% of the objects in the coded sentences.

Analyses

Each verb was assigned a relative frequency score on each variable: Transitivity was calculated by dividing the number of transitive utterances for each verb by the total number of coded utterances for that verb. Overall animacy was calculated as the number of animate subjects divided by the total number of coded subjects. Intransitive subject animacy was the ratio of animate subjects in the intransitive frame to the total number of coded intransitive subjects. Transitive subject animacy was not analyzed separately because transitive subjects were almost uniformly animate. Object animacy was the ratio of animate objects to total coded objects.

Classification In their classification studies, M&S's machine-learning algorithm learned to classify the verbs via explicit feedback as to the proper classification of a training subset of verbs. Supervised learning procedures of this type are generally considered a poor model for ordinary language acquisition, as children receive no direct feedback about the proper classification of verbs they have learned. To approximate this feature of language acquisition, we chose to classify the verbs using k-means clustering, an algorithm that does not receive direct feedback about correct classification. The k-means algorithm takes as its input scores on p variables (i.e. transitivity, etc.) for n objects (i.e. verbs) and attempts to assign the objects to k clusters. Each object is treated as a point in a p-dimensional space with its location determined by its scores on each variable. Initial divisions of this space into clusters are iteratively reorganized to minimize the sum, over all clusters, of the within-cluster distance between each point and its cluster center. Since k-means clustering is sensitive to the initial partitioning of the data, for each analysis reported below the clustering procedure was repeated 100 times with random initial clusterings and the solution with the lowest final within-cluster distance was used. Separate k-means analyses were conducted for each variable, as well as one analysis combining all variables.

The k-means procedure requires that the number of clusters be specified in advance. In the analyses reported below, the algorithm always divided the verbs into 2 clusters. Pilot analyses using 3 clusters produced poor results. Obviously, children do not have *a priori* knowledge of the number of classes into which they should group these verbs. How children go about discovering the correct number of classes over diverse sets of verbs remains to be investigated in future analyses.

Cluster evaluation Two measures were used to evaluate the resulting clusters. For each clustering solution, an accuracy score (*Acc*) was calculated by first assigning to each cluster the class label of the majority of its members. A verb was considered correctly classified if its actual class matched the class label of the cluster in which it was placed. *Acc* was calculated as the number of verbs correctly classified divided by the total number of verbs.

Acc scores can be relatively high for a clustering solution if one cluster is fairly uniform, even if the other cluster is very poor. The second evaluation metric, the adjusted Rand index (R_{adj}), measures the overall quality of the clustering solution (Hubert & Arabie, 1985). This index measures the similarity between the clustering solution and the true classification by examining all pairwise verb comparisons and classifying them as either agreements or disagreements. For example, placing two unspecified-object verbs in the same cluster would be considered an agreement, while placing them in different clusters would be a disagreement. The adjusted Rand index is scaled such that 1 indicates perfect agreement between the true classification and the clustering solution, while a value near 0 represents a random grouping (negative values can occur for extremely poor clustering solutions).

To provide a baseline for *Acc* and R_{adj}, we performed k-means analyses on 5,000 random permutations of the data. *Acc* and R_{adj} were averaged across the 5,000 permutations to yield a mean baseline for each score. Since different subsets of variables might yield different baselines, this procedure

was repeated for all subsets used in the experimental analyses. Across all data sets, the mean *Acc* was .57 and the mean R_{adj} was 0.00. To assess whether a given clustering solution represented a significant departure from baseline, we used the set of scores resulting from the random permutations as a reference distribution. The *p*-values shown in Table 3 are the proportion of randomly generated R_{adj} scores that were as extreme as or more extreme than each obtained score.

Results

Table 2 shows the relative frequency scores for each variable, separately by verb class. Causal verbs were less likely to have animate subjects than were unspecified-object verbs ($t(27)$=3.53, p<.005), as reported by M&S. The two classes of verbs also differed strikingly in the animacy of intransitive subjects: over half of the intransitive utterances containing causal verbs had inanimate subjects, while those containing unspecified-object verbs almost always had animate subjects ($t(27)$=7.15, p<.001). The two classes differed only marginally in object animacy ($t(27)$=1.91, p=.07). In contrast to M&S's findings, causal and unspecified-object verbs did not differ in transitivity in this sample.

Results of the experimental cluster analyses appear in Table 3. The clustering solution based on all 4 variables yielded a remarkably accurate classification, grouping 26 of the 29 verbs correctly. Its R_{adj} of .62 represents a significant improvement over baseline.

Individual feature analyses were conducted to assess the importance of each variable to the classification. As expected from the frequency scores, transitivity contributed little to the classification, yielding an essentially random grouping of the verbs. Classifications based on either overall subject animacy or object animacy resulted in *Acc* scores of only .62. While this is a 5% improvement over baseline, the *p*-values for the R_{adj} scores indicate that this improvement is not significant. The clustering based on intransitive subject animacy performed much better, correctly classifying 25 of the 29 verbs, just one verb fewer than the full analysis.

Table 2: Mean (SD) relative frequency scores by verb class

	Causal Verbs	Unspecified-Object Verbs
Transitivity	.70 (.20)	.67 (.21)
Subject Animacy	.86 (.12)	.98 (.05)
Intransitive Subject Animacy	.44 (.25)	.96 (.09)
Object Animacy	.15 (.20)	.34 (.32)

Table 3: Clustering results

Data set	*Acc*	R_{adj}	*p*-value
Transitivity	.55	-.02	.71
Animacy	.62	.04	.18
Intransitive Animacy	.86	.51	.0002
Object Animacy	.62	.03	.21
All variables	.90	.62	.0000

With the exception of the full analysis, no other possible combination of the variables yielded a result better than the intransitive subject animacy classification.

Discussion

The results of this study demonstrate that surface features can be used to distinguish causal from unspecified-object verbs in child-directed speech with high accuracy. In addition, the individual feature analyses, while they deviated in their details from those reported by M&S, revealed an important underlying similarity in the nature of the surface features that were informative in very different corpora.

The differences first: M&S predicted and found that transitivity was useful in the classification of unspecified-object and causal verbs. Causal verbs occurred less often in the transitive frame. M&S predicted this outcome on markedness grounds: Since causal alternation verbs encode in their transitive frame complex events consisting of action and result subevents, while unspecified-object verbs do not, they argued that the causal verbs should be used transitively less often. We found no such tendency in our data. In our corpora, both causal and unspecified-object verbs demonstrated equally high rates of transitive use. As a result, transitivity was not useful in classifying the verbs in our data set.

This disparity likely results from differences between the WSJ and our CHILDES corpora. As mentioned previously, Roland and Jurafsky (1998) have found that subcategorization differences across corpora can be partially attributed to differences in the distribution of verb senses within those corpora. For instance, in our data the verb *fold* was used primarily in the context of doing laundry (e.g. "Let's fold it [the towel] nice and neat.") and therefore the uses of this verb were predominantly transitive (69%). M&S reported that *fold* was used transitively only 23% of the time in their WSJ data. It is unlikely that the WSJ contains the sense of *fold* that pertains to laundry. Instead, in the WSJ *fold* is probably used to refer to things like the collapse of a corporation (e.g. "After serious financial trouble, the company folded."). Although we did not code for verb sense, it seems likely that differences in discourse context and style between the two corpora led to many sense differences like those seen with *fold* and that these different senses gave rise to the different transitivity patterns.

M&S also found that the animacy of the subject noun phrase aided in classifying the two groups of verbs. Due to inherent differences in argument structure (see above), the causal verbs displayed lower subject animacy than the unspecified-object verbs. Although this difference was also present in our data, it was not sufficiently large to classify the verbs properly. The high rate of transitivity in our data likely contributed to this difference between our findings and M&S's. While there is a principled reason for these verbs to differ in the animacy of their intransitive subjects, the same is not true of the transitive subjects: Both groups assign an agent to the subject of the transitive, so transitive subjects should be largely animate, regardless of verb class. This prediction was borne out in our data. Since there was a higher percentage of

transitive utterances in our data set than in M&S's, any difference between the groups was obscured by the uniformly animate transitive subjects.

Once the transitive subjects were removed, the difference between the two groups became apparent. Intransitive subject animacy classified most of the verbs correctly, performing much better than overall subject animacy did. Thus, animacy information proves to be a robust indicator of verb class, holding up across two very different sets of corpora. While our analysis required a more sensitive measure, intransitive subject animacy, the conclusion remains as Merlo and Stevenson predicted: Surface patterns of animacy reflect the underlying argument structure of these verb classes.

Implications for verb learning

Our results suggest that animacy information could provide children with a highly informative cue about a verb's class. Can children track animacy information of this type? Many researchers have documented that subjects are preferentially animate (e.g. Bock, Loebell, & Morey, 1992) and that children are sensitive to this tendency from very early in life. Two- and 4-year-olds more easily comprehend transitive sentences with animate than with inanimate subjects (Corrigan, 1988). A similar bias is seen in production, where animate subjects facilitate children's use of the passive construction (Lempert, 1984). These studies demonstrate that children are sensitive to animacy information present in the input and they expect that subjects should be animate. Given this expectation, the occurrence of an inanimate entity in subject position might be particularly salient.

Recent work also shows that children store semantically-laden combinatorial information for individual verbs, and retrieve that information when they identify a verb in the input. When children hear a semantically restrictive verb (e.g. *eat*), they quickly locate a potential direct object of that verb (Chang & Fernald, 2003). Similarly, 2- and 2.5-year-olds used knowledge of the semantic restrictions of a verb to learn a novel noun (Goodman, McDonough, & Brown, 1998). When presented with an array of objects and the sentence "Mommy feeds the *ferret*," children correctly assumed that *ferret* referred to the only animate entity present. Taken together with the evidence discussed above, these findings suggest that children can track animacy information, and more specific semantic restrictions, about particular arguments of individual verbs.

The problem of verb sense Grimshaw (1994) points out that although there is a principled relation between a particular verb sense and the set of frames associated with that sense, there is no clear relationship between a verb stem (combining all its senses) and its full set of subcategorization frames. As a result, subcategorization information provides useful cues to verb meaning only if it is tracked separately for each sense of a verb. Grimshaw therefore predicts that learning that depends on subcategorization information will be highly errorful.

Our results suggest that this is not necessarily the case, given the statistics of child-directed speech. The coding procedures conflated all senses of the verb into a single representation. On Grimshaw's account, collapsing across senses in this fashion should have produced very poor learning, yet our analysis successfully classified a large majority of the verbs. Further inspection of the data reveals that for many verbs, one sense was very frequent, dominating all other senses (e.g., *fold*, see above). For other verbs, such as *change*, several senses of the verb occurred throughout the corpora (e.g. "It [the picture] changes when you move it." vs. "Shall I change your diaper?"), but the subcategorization profiles across senses were complementary rather than contradictory (i.e. one sense occurred primarily in the intransitive, while the other occurred in the transitive). Both of these patterns resulted in a successful classification.

The three verbs that were never correctly classified (*move, slide, turn*) appeared to have different senses that behaved differently. One such verb, *turn*, was frequently used in an activity sense (e.g. "You're turning!" meaning "turning yourself"). Since this intransitive use occurred with an animate subject, *turn* was classified with the unspecified-object verbs rather than the causal verbs in all analyses. Though this analysis correctly grouped these activity uses of *turn* with other activity verbs in the unspecified-object class, it failed to reflect the fact that *turn* also has a causal sense. The fact that this only posed a problem for 3 of the 29 verbs, however, suggests that children could use subcategorization information to learn the meanings of many common verbs without requiring prior knowledge of verb sense.

Acknowledgments

References

Bard, E. G., & Anderson, A. H. (1994). The unintelligibility of speech to children: Effects of referent availability. *Journal of Child Language, 21,* 623-648.

Bloom, L., Hood, L., & Lightbown, P. (1974). Imitation in language development: If, when and why. *Cognitive Psychology, 6,* 380–420.

Bloom, L., Lightbown, P., & Hood, L. (1975). Structure and variation in child language. *Monographs of the Society for Research in Child Development, 40,* (Serial No. 160).

Bock, J. K., Loebell, H., & Morey, R. (1992). From conceptual roles to structural relations: Bridging the syntactic cleft. *Psychological Review, 99,* 150-171.

Brown, R. (1973). *A first language: The early stages*. Cambridge, MA: Harvard University Press.

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*.

Childers, J. B., & Tomasello, M. (2001). The role of pronouns in young children's acquisition of the English transitive construction. *Developmental Psychology, 37,* 739-748.

Clark, E. V. (1982). The young word maker: A case study of innovation in the child's lexicon. In E. Wanner & L. R. Gleitman (Eds.), *Language acquisition: The state of the*

art. Cambridge, MA: Cambridge University Press.

Corrigan, R. (1988). Children's identification of actors and patients in prototypical and nonprototypical sentence types. *Cognitive Development, 3,* 285-297.

Demetras, M. (1989). *Working parents' conversational responses to their two-year-old sons.* University of Arizona.

Fisher, C. (2000). From form to meaning: A role for structural analogy in the acquisition of language. In H. W. Reese (Ed.), *Advances in child development and behavior.* New York: Academic Press.

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.

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

Gleitman, L. R. (1990). The structural sources of verb meanings. *Language Acquisition, 1,* 3-55.

Gleitman, L. R., Cassidy, K., Nappa, R., Papafraqou, A., & Trueswell, J. C. (2005). Hard Words. *Language Learning and Development, 1,* 23-64.

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.

Grimshaw, J. (1990). *Argument structure.* Cambridge, MA: MIT Press.

Grimshaw, J. (1994). Lexical reconciliation. *Lingua, 92,* 411-430.

Higginson, R. P. (1985). *Fixing-assimilation in language acquisition.* Unpublished doctoral dissertation. Washington State University.

Hubert, L., & Arabie, P. (1985). Comparing partitions. *Journal of Classification, 2,* 193-218.

Jackendoff, R. (1990). *Semantic structures.* Cambridge, MA: MIT Press.

Kuczaj, S. A. (1986). General developmental patterns and individual differences in the acquisition of copula and auxiliary be forms. *First Language, 6,* 111–117.

Laakso, A., & Smith, L. (2004). Pronouns predict verb meanings in child-directed speech. In K. Forbus, D. Gentner & T. Regier (eds.), *Proceedings of the 26th Annual Meeting of the Cognitive Science Society.* Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

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

Lempert, H. (1984). Topic as a starting point for syntax. *Monographs of the Society for Research in Child Development, 49.*

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

Levin, B., & Rappaport-Hovav, M. (2005). *Argument realization.* New York: Cambridge University 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.

Mintz, T. H. (2003). Frequent frames as a cue for grammatical categories in child-directed speech. *Cognition, 90,* 91-117.

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. & Kako, E. (1993). First contact in verb acquisition: Defining a role for syntax. *Child Development, 64,* 1665-1687.

Newport, E. L., Gleitman, H., & Gleitman, L. R. (1977). Mother, I'd rather do it myself: Some effects and noneffects of maternal speech style. In C. E. Snow, & C. A. Ferguson, *Talking to children: language input and acquisition.* Cambridge: Cambridge University Press.

Ninio, A., Snow, C., Pan, B., & Rollins, P. (1994). Classifying communicative acts in children's interactions. *Journal of Communications Disorders, 27,* 157-188.

Pinker, S. (1989). *Learnability and cognition.* Cambridge, MA: MIT Press.

Post, K. (1994). Negative evidence. In J. Sokolov & C. Snow (Eds.), *Handbook of Research in Language Development Using CHILDES.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Roland, D., & Jurafsky, D. (1998). *Computing verbal valence frequencies: corpora versus norming studies.* Poster session presented at the CUNY sentence processing conference, Santa Monica, CA.

Suppes, P. (1974). The semantics of children's language. *American Psychologist, 29,* 103–114.

Warren-Leubecker, A. (1982). *Sex differences in speech to children.* Unpublished doctoral dissertation. Georgia Institute of Technology.

Yuan, S., & Fisher, C. (in press). "Really, he blicked the cat?": 2-year-olds learn distributional facts about verbs in the absence of a referential context. *Proceedings of the 29th Annual Boston University Conference on Language Development.* Boston, MA: Cascadilla Press.