

3 Discovering Verbs Through Multiple-Cue Integration

Morten H. Christiansen and Padraic Monaghan

Introduction

Before children can ride a bicycle or tie their shoes, they have learned a great deal about how words are combined to form complex sentences. This achievement is especially impressive because children acquire most of this syntactic knowledge with little or no direct instruction. Nevertheless, mastering natural language syntax may be among the most difficult learning tasks that children face. In adulthood, syntactic knowledge can be characterized by constraints governing the relationship between grammatical categories of words (such as noun and verb) in a sentence. However, acquiring this knowledge presents the child with a chicken-and-egg problem: the syntactic constraints presuppose the existence of grammatical categories because syntactic knowledge is generally couched in terms of categories of words and not in terms of individual words. On the other hand, grammatical categories have little value in and by themselves; rather, they are only useful insofar as they support syntactic constraints. A similar "bootstrapping" problem faces a student learning an academic subject such as physics: understanding momentum or force presupposes some understanding of the physical laws in which they figure, yet these laws presuppose the very concepts they interrelate. But the bootstrapping problem solved by young children seems vastly more challenging, both because the constraints governing natural language are so intricate and because young children do not have the intellectual capacity or explicit instruction available to the academic student. So how does the child solve the bootstrapping problem in language acquisition? In this chapter, we pursue a possible solution in the form of multiple-cue integration.

By 1 year, infants will have learned a great deal about the sound structure of their native language (for reviews, see Jusczyk, 1997, 1999; Kuhl, 1999; Pallier, Christophe, & Mehler, 1997; Werker & Tees, 1999). Thus, when they face the

problem of bootstrapping syntax at the beginning of their second year, they are already well acquainted with the phonological and prosodic regularities of their native language. The multiple-cue integration hypothesis suggests that this perceptual attunement provides an essential scaffolding for later learning by biasing children toward aspects of the input that are particularly informative for acquiring syntactic information (e.g., Christiansen & Dale, 2001; Gleitman & Wanner, 1982; and contributions in Morgan & Demuth, 1996; Weissenborn & Höhle, 2001). Specifically, the integration of multiple probabilistic cues derived from the co-occurrence of words (distributional), their sound properties (phonological), and their intonational (prosodic) as well as situational (semantic) context by perceptually attuned general-purpose learning mechanisms may hold the key to how children solve the bootstrapping problem. In this way, multiple cues can provide reliable evidence about linguistic structure that is unavailable from any single source of information.

A further initial strategy that a child may bring to bear on the problem of bootstrapping syntax is to focus on the discovery of nouns and verbs, perhaps the most salient groups of content words in that they refer to objects and actions in the environment. Even in this much reduced version, children's learning task remains formidable given that they are still in the process of making sense of the nonlinguistic world as well. And it is in this context that verbs may be particularly difficult to pin down. Minimally, early verb learning requires that children master three different complex learning tasks. First, children need to be able to segment fluent speech to locate possible verb forms using distributional, acoustic, and other types of language-internal cues (see the other chapters in part I). Second, they must be able to find the appropriate parts of actions to be named among unfolding event sequences involving many types of language-external cues (see chapters in part II). Finally, they have to learn to integrate language-internal and language-external cues in the service of acquiring the form and meaning of verbs (and other words; see chapters in part III).

In this chapter, we discuss how children may accomplish the difficult task of verb learning, focusing on the integration of multiple language-internal cues to verb forms. We first review previous work on multiple-cue integration. We then report results from novel analyses of corpora of English child-directed speech, pointing to different roles for distributional and phonological cues in the learning of nouns and verbs. Finally, we relate the differential roles of cues to differences in semantic support for nouns and verbs in language-external information, and discuss possible implications of our results for the understanding of word learning more generally.

Multiple-Cue Integration in Language Acquisition

There are three sources of information that children could potentially bring to bear on solving the bootstrapping problem: innate knowledge in the form of

linguistic universals; language-external information concerning observed semantic relationships between language and the world; and language-internal information, such as aspects of phonological, prosodic, and distributional patterns that indicate the relation of various parts of language to each other.

Although some kind of innate knowledge may play a role in language acquisition, it cannot solve the bootstrapping problem. Even with built-in abstract knowledge about grammatical categories and syntactic rules (e.g., Pinker, 1984), the bootstrapping problem remains formidable: Innate knowledge can only help address the bootstrapping problem by building in universal aspects of language, and relationships between words and grammatical categories clearly differ between languages (e.g., the sound /su/ is a noun in French, *sou*, but a verb in English, *sue*). Crucially, children still have to map the right sound strings onto the right grammatical categories while determining the specific syntactic relations between these categories in their native language. Moreover, there now exists strong experimental evidence that children do not initially use abstract linguistic categories but instead employ novel words as concrete items, thereby challenging the usefulness of hypothesized innate grammatical categories (Tomasello, 2000). Thus, independently of whether or not innate linguistic knowledge is hypothesized to play an important role in language acquisition, it seems clear that other sources of information nevertheless are necessary to solve the bootstrapping problem.

Language-external information is likely to contribute substantially to language acquisition. Correlations between environmental observations relating prior semantic categories (e.g., objects and actions) and grammatical categories (e.g., nouns and verbs) may furnish a "semantic bootstrapping" solution (Pinker, 1984). However, given that children acquire linguistic distinctions with no semantic basis (e.g., gender in French; Karmiloff-Smith, 1979), semantics cannot be the only source of information involved in solving the bootstrapping problem. Another extralinguistic factor is cultural learning, whereby children may imitate the pairing of linguistic forms and their conventional communicative functions (Tomasello, Kruger, & Ratner, 1993). For example, by observing the idiom "John spilled the beans" used in the appropriate context, the child by reproducing it can discover that it means that John has revealed some sort of secret and not that he is a messy eater. However, to break down the linguistic forms into relevant units, it appears that cultural learning must be coupled with language-internal learning.

Though not the only source of information involved in language acquisition, we suggest that language-internal information is fundamental to bootstrapping the child into syntax. However, although language-internal input appears to be rich in potential cues to linguistic structure, there is an important caveat: the individual cues are only partially reliable, and none considered alone provides an infallible bootstrap into language. Thus, a learner could use the tendency for English nouns to be longer than verbs to determine that *elephant* is a noun, but the same strategy would fail for *investigate*. Likewise, although speakers tend to pause at linguistically meaningful places in a sentence (e.g., following a phrase or a clause), pauses

also occur elsewhere. And although it is a good distributional bet that a determiner (e.g., *the*) will be followed by a noun, there are other possibilities (e.g., adjectives, such as *big*). To acquire language successfully, it seems that the child needs to integrate a great diversity of multiple probabilistic cues to language structure in an effective way. Fortunately, as we shall see next, there is a growing bulk of evidence showing that multiple probabilistic cues are available in language-internal input, that children are sensitive to them, and that they facilitate learning through multiple-cue integration.

Bootstrapping Through Multiple Language-Internal Cues

We distinguish between three types of language-internal cues: phonological, prosodic, and distributional cues. Phonological information—including stress, vowel quality, and duration—may help distinguish grammatical function words (e.g., determiners, prepositions, and conjunctions) from content words (nouns, verbs, adjectives, and adverbs) in English (e.g., Cutler, 1993; Gleitman & Wanner, 1982; Monaghan, Chater, & Christiansen, 2005; Morgan, Shi, & Allopenna, 1996; Shi, Morgan, & Allopenna, 1998). Phonological information may also help distinguish between nouns and verbs (Monaghan et al., 2005). For example, adults are sensitive to the fact that English disyllabic nouns tend to receive initial-syllable (trochaic) stress whereas disyllabic verbs tend to receive final-syllable (iambic) stress (Kelly, 1988). Moreover, acoustic analyses have shown that even noun-verb ambiguous disyllabic words that change grammatical category but not stress placement can be differentiated by syllable duration and amplitude cue differences (Sereno & Jongman, 1995). Experiments indicate that children as young as 3 years old are sensitive to this stress cue, even though few multisyllabic verbs occur in child-directed speech (Cassidy & Kelly, 1991, 2001). Other potential noun-verb cues in English include differences in word duration, consonant voicing, and vowel types—many of these cues may also be relevant cross-linguistically (see Kelly, 1992, for a review).

Prosodic information provides cues for word and phrasal/clausal segmentation and may help uncover syntactic structure (e.g., Gerken, Jusczyk, & Mandel, 1994; Gleitman & Wanner, 1982; Kemler Nelson, Hirsh-Pasek, Jusczyk, & Wright Cassidy, 1989; Morgan, 1996). Acoustic analyses suggest that differences in pause length, vowel duration, and pitch indicate phrase boundaries in both English and Japanese child-directed speech (Fisher & Tokura, 1996). Infants seem highly sensitive to such language-specific prosodic patterns (Gerken et al., 1994; Kemler-Nelson et al., 1989; for reviews, see Gerken, 1996; Jusczyk & Kemler-Nelson, 1996; Morgan, 1996)—a sensitivity that may start in utero (Mehler et al., 1988). Prosodic information also improves sentence comprehension in 2-year-olds (Shady & Gerken, 1999). Results from artificial language learning experiments with adults show that prosodic marking of syntactic phrase boundaries facilitates learning (Morgan, Meier, & Newport, 1987; Valian & Levitt, 1996). Evidence from event-related

brainwave potentials in adults showing that prosodic information has an immediate effect on syntactic processing (Steinhauer, Alter, & Friederici, 1999) further underscores the importance of this cue. Unfortunately, prosody is also partly affected by a number of nonsyntactic factors such as breathing patterns, resulting in an imperfect mapping between prosody and syntax (Fernald & McRoberts, 1996). Nonetheless, infants' sensitivity to prosody provides a rich potential source of syntactic information (Fisher & Tokura, 1996; Gerken 1996; Morgan, 1996).

Information about the distribution of linguistic fragments at or below the word level may also provide cues to grammatical category. Morphological patterns across words may be informative—for example, English words that are observed to have both *-ed* and *-s* endings are likely to be verbs (Maratsos & Chalkley, 1980). Artificial language learning results show that adults are better at learning grammatical categories cued by word internal patterns (Brooks, Braine, Catalano, & Brody, 1993; Frigo & McDonald, 1998). Corpus analyses have demonstrated that distributional patterns of word co-occurrence also give useful cross-linguistic cues to grammatical categories in child-directed speech (e.g., Mintz, 2003; Monaghan et al., 2005; Redington, Chater, & Finch, 1998; Redington et al., 1995). Given that function words primarily occur at phrase boundaries (e.g., initially in English and French, finally in Japanese) they may reveal syntactic structure. This is confirmed by corpus analyses (Mintz, Newport, & Bever, 2002) and results from artificial language learning (Green, 1979; Morgan et al., 1987; Valian & Coulson, 1988). Finally, artificial language learning experiments indicate that duplication of morphological patterns across phrase-related items (e.g., Spanish: *Los Estados Unidos*) facilitates learning (Meier & Bower, 1986; Morgan et al., 1987).

Phonological information may help distinguish between function and content words and between nouns and verbs. Prosodic information provides cues for word and phrasal or clausal segmentation and may help uncover syntactic structure. Distributional information affords cues for labeling and segmentation and perhaps evidence towards syntactic relations. None of these cues in isolation suffice to solve the bootstrapping problem; rather, they must be integrated to overcome the limited reliability of individual cues. Recent connectionist simulations have demonstrated that efficient and robust learning mechanisms exist for multiple-cue integration (Christiansen & Dale, 2001; Reali, Christiansen, & Monaghan, 2003). Despite previous theoretical reservations about the value of multiple-cue integration (Fernald & McRoberts, 1996), analyses of network performance revealed that learning under multiple cues results in faster, better, and more uniform learning. Moreover, the networks were able to distinguish between relevant cues and distracting cues, and performance did not differ from networks that received only reliable cues. The efficacy of multiple-cue integration has also been confirmed in artificial language learning experiments (Billman, 1989; Brooks et al., 1993; McDonald & Plauche, 1995; Morgan et al., 1987).

After one year of exposure to spoken language, children's perceptual attunement is likely to allow them to utilize language-internal probabilistic cues (for reviews, see Jusczyk, 1997, 1999; Kuhl, 1999; Pallier et al., 1997; Werker & Tees, 1999). For example, infants appear sensitive to the acoustic differences between function and content words (Shi, Werker, & Morgan, 1999) and the relationship between function words and prosody in speech (Shafer, Shucard, Shucard, & Gerken, 1998). Young infants can detect differences in syllable number among isolated words (Bijeljac, Bertocini, & Mehler, 1993)—a possible cue to noun-verb differences. Moreover, infants are accomplished distributional learners (e.g., Gómez & Gerken, 1999; Saffran, Aslin, & Newport, 1996; see Gómez & Gerken, 2000; Saffran, 2003, for reviews), and importantly, they are capable of multiple-cue integration (Mattys, Jusczyk, Luce, & Morgan, 1999; Morgan & Saffran, 1995). When solving the bootstrapping problem, children are also likely to benefit from specific properties of child-directed speech, such as the predominance of short sentences (Newport, Gleitman, & Gleitman, 1977) and exaggerated prosody (Kuhl et al., 1997).

This review has indicated that a range of language-internal cues are available for language acquisition, that these cues affect learning and processing, and that mechanisms exist for multiple-cue integration. Next we present two sets of corpus analysis experiments showing that language-internal cues appear to take on different roles in the context of learning about nouns and verbs in English. Specifically, we have found that phonological cues may be particularly important for learning verbs, whereas distributional information may be more useful for learning about nouns.

Experiment 1: The Importance of Phonological Cues for Verb Learning

In a previous study (Monaghan et al., 2005), we quantified the potential usefulness of phonological and distributional cues for distinguishing between nouns and verbs through a series of corpus analyses of a large corpus of child-directed speech. The method we used to assess this usefulness was discriminant analysis, which attempts to use the cues to carve up a set of words into distinct categories. The extent to which this can be done effectively, matching the actual syntactic categories, can then be assessed. Figure 3.1 is a schematic diagram of the method of distinguishing nouns and verbs, which are represented as dots in a space determined by the set of cues. In essence, discriminant analysis provides a hyperplane through the word space, depicted by the gray-shaded surface, based on the cues that most accurately reflect the actual category distinction. In the figure, the discriminant analysis classifies nouns and verbs effectively, with most nouns occurring above the hyperplane and most verbs positioned below the plane. Our preferred method is to use a "leave-one-out cross-validation" method, which is

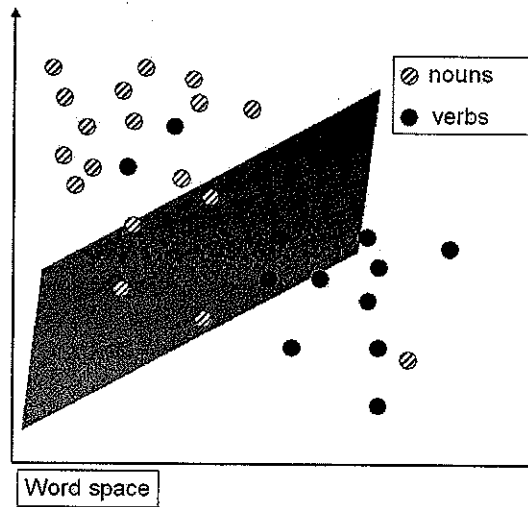


Figure 3.1. Schematic illustration of discriminant analysis of nouns and verbs.

a conservative measure of classification accuracy and works by assessing the accuracy of the classification of words that are not used in positioning the hyperplane. This means that the hyperplane is constructed on the basis of the information on all words except one, and then the classification of the omitted word is assessed. This is then repeated for each word, and the overall classification accuracy can then be determined. The results of the analyses of phonological and distributional cues showed that the use of several cues provides not only more accurate classification than single cues, but better generalization to novel situations. We also found that frequency interacted with the two cues: Distributional cues were more useful for high-frequency words, whereas phonological cues were more reliable for low-frequency words. The integration of phonological and distributional cues resulted in 66.7% correct noun-verb classification. These results indicated the general usefulness of multiple-cue integration in learning syntactic categories. In Experiment 1, we extend this work to determine whether phonological and distributional cues may be differentially useful for learning nouns and verbs. In this study, we generated phonological and distributional cues for the 1,000 most frequent words in child-directed speech and assessed the extent to which these different cue types distinguished each syntactic category from all others.

Method

Corpus Preparation

In order to effectively model the language environment of the child, we used a corpus of child-directed speech. We concatenated all the child-directed speech

from the English CHILDES database (MacWhinney, 2000). This amounted to 5,436,855 words. All pauses and turn-taking were marked as utterance boundaries, resulting in approximately 1.4 million utterances. For each word, we derived its phonological form from the CELEX database (Baayen, Pipenbrock, & Gulikers, 1995). Orthographic forms with alternative pronunciations were assigned the most frequent pronunciation from CELEX. Many words in English are syntactically ambiguous: *chair* and *table* have usage as both nouns and verbs, though in most cases there is a clear most frequent usage (*chair* and *table* are usually nouns). CELEX records the frequency of usage of words as different syntactic categories, and so we used this information to label each word with its most frequent usage. The syntactic categories distinguished in CELEX were nouns, adjectives, numerals, verbs, articles, pronouns, adverbs, prepositions, conjunctions, interjections, and contractions (such as *should've*, *can't*). We counted the number of occurrences of each word in the child-directed speech corpus and selected the 1,000 most frequent words from the CHILDES database and hand-coded those words that did not appear in the CELEX database (mainly comprising proper nouns, interjections, and alternative spellings, such as *mama*, *mamma*).

Phonological Information

We coded each word for 16 phonological cues that have been posited as useful in distinguishing different syntactic categories in English (Campbell & Besner, 1981; Cutler, 1993; Cutler & Carter, 1987; Gleitman & Wanner, 1982; Kelly, 1992; Kelly & Bock, 1988; Marchand, 1969; Morgan et al., 1996; Shi et al., 1996; Sereno & Jongman, 1990). These cues operate either at the word level, the syllable level, or the phoneme level. Word-level cues concern the properties of the word taken as a whole; syllable-level cues assess the distribution of consonants and phonemes in the syllable; and phoneme-level cues assess the properties of particular phonemes within the word. Table 3.1 summarizes the cues we used in the analyses; for more details of the precise encoding of the cues, see Monaghan et al. (2005).

Distributional Information

The child-directed speech corpus provided the contexts within which each word occurred in one or more utterances, and our distributional cues were designed to exploit this information. We examined the local context of each word from CHILDES in order to gain a reflection of the value of distributional information for categorizing each word. As our measure of distributional information we selected a set of high-frequency words and determined the number of times each word occurred immediately succeeding each of these high-frequency words. The rationale for this was that high-frequency words can be used to promote the categorization of lower-frequency words that follow them. Valian and Coulson (1988) constructed an artificial language in which low-frequency nonsense words

Table 3.1 Phonological cues used in the discriminant analyses of syntactic category

Phonological Cue	Examples		
	<i>monkey</i>	<i>stretched</i>	<i>that</i>
Word level			
Length in phonemes	5	6	3
Length in syllables	2	1	1
Presence of stress	1	1	0
Syllable position of stress	1	1	0
Syllable level			
Number of consonants in word onset	1	3	1
Proportion of phonemes that are consonants	0.6	0.83	0.67
Proportion of syllables containing reduced vowel	0	0	1
Reduced 1st vowel	0	0	1
-ed inflection	0	1	0
Phoneme level			
Proportion of consonants that are coronal	0	1	1
Initial /Δ/	0	0	1
Final voicing	0	2	2
Proportion of consonants that are nasals	0.67	0	0
Position of stressed vowel (1: front, 3: back)	3	1	2*
Position of vowels (1: front, 3: back)	2	1	1
Height of vowels (0: high, 3: low)	1.25	2	2.5

*Words with no stress were assigned the median position.

followed only one of two high-frequency nonsense words. Participants were able to learn to group together the lower-frequency words that followed each of the high-frequency words. We selected the 20 most frequent words from the CHILDES database and counted the times that these context words preceded the target word. We then calculated the signed log-likelihood test score for each context-target word pairing (Dunning, 1993; Monaghan et al., 2005). This score reflects the extent to which the target word co-occurs more or less than by chance with the context word in the corpus. High positive values indicate there is a greater than chance co-occurrence between the words, whereas large negative values reflect that the co-occurrence occurs less than by chance. Values close to zero indicate that co-occurrence occurs at chance level. Each of the 1,000 words we examined, therefore, had log-likelihood test scores for the 20 most frequent context words.

Analysis

The potential value of cues for distinguishing the syntactic categories of words was assessed using leave-one-out cross-validation discriminant analysis. For each

syntactic category, we labeled words as either belonging to that category or belonging to another category. We used stepwise discriminant analysis which meant that cues were only entered into the analysis if they contributed significantly to making the distinction between words in the category in question and all other words. We performed discriminant analyses employing either the phonological or the distributional cues. In order to compare classification performance for the different syntactic categories, successful classification was weighted by the size of the category—so accuracy greater than 50% reflected better-than-chance performance.

Results and Discussion

Table 3.2 indicates the correct classification of words for each category distinction in the discriminant analyses, employing either phonological or distributional cues. It was not the case that any type of cue resulted in perfect classifications of any syntactic category, though classification was better than chance level of 50% for all analyses. The results supported the prediction that using phonological cues alone resulted in better classification of verbs than of nouns. The results also supported the claim that using distributional cues alone would provide better classification of nouns. Also of note in the results are the particular usefulness of phonological cues for classifying function words—pronouns, prepositions, conjunctions, and interjections all had more than 80% correct classifications with these cues, whereas distributional cues were poor for these word categories.

In order to determine in more detail the role of phonological and distributional cues for distinguishing nouns and verbs, we examined which words were correctly classified by both the phonological and distributional analysis and

Table 3.2 Discriminant analysis classification results using phonological cues and distributional cues for distinguishing each category from all other words

Category	Phonological Classification	Distributional Classification
Noun	61.8	64.3
Adjective	54.8	59.7
Numeral	55.7	62.3
Verb	66.4	62.3
Article	46.9	49.9
Pronoun	83.6	62.9
Adverb	68.1	52.4
Preposition	85.3	62.3
Conjunction	85.6	53.2
Interjection	86.7	57.2
Contraction	63.7	56.4

explored which words were correctly classified using phonological cues only but which were incorrectly classified using distributional cues, and vice versa, the case where distributional cues produced a correct classification but phonological cues resulted in incorrect classification. Finally, we also encoded those words that were classified incorrectly by both analyses. There were two possibilities for the resulting classifications. It may be that the same words are correctly classified by analyses based on both cue types, or it may be that there is complementarity in the classifications: Those words incorrectly classified by, say, the phonological cues, may be correctly classified by the distributional cues. Table 3.3 presents the results, showing the number of words on which the classifications agreed and disagreed. Hierarchical log-linear analysis can then be used to assess whether there are main effects and interactions between the classifications based on the different cue types and the noun-verb category. Table 3.4 shows the results of the one-, two-, and three-way log-linear analyses on the table shown in table 3.3. The one-way analyses refer to main effects in the table, the two-way analyses refer to interactions between two of the factors, and the three-way analysis tests whether there is a three-way interaction in the table.

The one-way effect of noun-verb category can be explained by there being more nouns than verbs. The one-way effects of phonological cues and distributional cues reflected the fact that each classification assigned words to the correct category significantly more than by chance. The two-way effects of noun-verb by phonological cues and noun-verb by distributional cues indicate that the classifications were more successful overall for nouns than for verbs. The two-way effect of distributional cues by phonological cues was due to phonological cues being more

Table 3.3 Correct and incorrect classifications of nouns from other words and verbs from other words in the phonological and distributional discriminant analyses

Nouns		Distributional Classification		
Phonological Classification	Correct	Incorrect	Total	
Correct	405	182	592	
Incorrect	222	177	399	
Total	627	359	986	
Verbs		Distributional Classification		
Phonological Classification	Correct	Incorrect	Total	
Correct	406	155	561	
Incorrect	44	381	425	
Total	450	536	986	

Table 3.4 Results of hierarchical loglinear analysis

Generating Class	Likelihood Ratio χ^2	df
<i>One-way</i>		
Noun/verb	489.730	6
Phon	530.069	
Dist	551.201	
<i>Two-way</i>		
Noun/verb \times Phon	489.482	5
Noun/verb \times Dist	529.820	
Phon \times Dist	468.349	
<i>Three-way</i>		
Noun/verb \times Phon \times Dist	468.101	4

"Phon" refers to the classification resulting from the phonological cues, and "Dist" refers to the distributional cue classification. For all likelihood ratio χ^2 , $p < .001$.

effective in classifying words than the distributional cues. However, interpretation of these lower-level interactions must be moderated by the three-way interaction.

The three-way interaction suggests that the combination of phonological and distributional information operates differently for nouns and for verbs. The principal differences in the classifications in table 3.3 are the words that the phonological information classifies wrongly. For nouns, the distributional information operates to classify over 50% of these correctly, whereas for verbs, only 44 of 425 incorrect classifications are remedied by the distributional information. Whereas noun classification is benefited by both phonological and distributional information, the addition of distributional information for the verb classifications does not greatly improve classification based on the phonological information alone.

The results of assessing the effects of phonological and distributional information for discriminating words of each syntactic category fit our initial hypothesis well. Phonological information appears to be particularly useful for determining the verb category, but less useful for noun classifications, once distributional information has been taken into account. As an aside, the high accuracy of phonological cues for function words, shown in table 3.2, appears to indicate that phonological information is particularly useful and rich in all cases where distributional information is less accurate as a basis for classification.

Experiment 2: Distributional Cues Work Better for Nouns

The results of Experiment 1 could potentially depend on using a particular type of distributional information. In our Experiment 2, we therefore sought to determine

whether other approaches to distributional information may be more useful for distinguishing nouns and verbs as syntactic categories. We considered two forms of distributional information—one in which the information is only about the preceding word (Monaghan & Christiansen, 2004), and one, developed by Mintz (2003), which considers jointly the preceding and succeeding word. We predict that enriched distributional information will benefit classification of both nouns and verbs, but in both cases, classification of nouns will exceed that of verbs.

Mintz (2003) developed an intriguing approach to test the potential assistance that distributional information may provide in classifying words from different syntactic categories. He proposed that a frame formed by linking the preceding and succeeding word would predict with a high degree of accuracy the syntactic category of the intervening word. Put another way, given a particular pairing of preceding (*A*) and succeeding (*B*) words, all words (*x*) that occur within that pair (as *AxB*) will be of the same syntactic category. To test this prediction, Mintz selected the 45 most frequent *A_B* pairs of words from subsets of the CHILDES database. He then clustered together all words that occurred inside the *A_B* frame (i.e., *x* in the *AxB* frame). The resulting clusters tended to be of the same syntactic category with very high accuracy (see also chapter 1, this volume). Experiment 2 replicates and extends this study to investigate differentials in classification of nouns and verbs.

Method

Corpus Preparation

We replicated Mintz's (2003) analysis on one of the corpora that he employed: speech spoken to a child aged between 0 and 2 years, 6 months (anne01a-anne23b; Theakston, Lieven, Pine, & Rowland, 2001). We replaced all pauses and turn-taking with utterance boundaries, resulting in 93,269 words distributed over 30,365 utterances. As with the analysis of the large corpus in Experiment 1, syntactic category was taken from CELEX, and all words that were not in CELEX were hand coded. We used a slightly different set of syntactic categories to Experiment 1, to make our results more comparable to the categories used by Mintz: hence we distinguished words into the categories of noun, adjective, numeral, verb, article, pronoun, adverb, conjunction, preposition, interjection, *wh*-word (e.g., *what*, *where*), and proper noun. Contractions were classified according to the syntactic category of their first element, so *you're* was classified as pronoun, *could've* was classified as verb, and *what's* as *wh*-word.

Analysis

We then selected the 45 most frequent *A_B* frames from the Theakston et al. (2001) corpus and clustered together all words that intervened in each frame. For

example, for the frame *a _____ on* the following words occurred and were subsequently clustered together: *puddle, suck, welly-boot, plaster, stamp, tap, wee, sandcastle, spoon, fish, tunnel, bib, video, wee-wee, blanket, nappy, walk, cow, hat, ride, car*. We assessed accuracy by counting the number of correct pairings of words of the same syntactic category (hits) in each cluster, and dividing this by the number of pairings of all words within each cluster (hits + false alarms). In the above example for the *a _____ on* frame, accuracy was 100% for this cluster. We assessed completeness by dividing the number of hits by the number of pairings of words of the same syntactic category that occurred in all clusterings (hits + misses). For the *a _____ on* frame, any nouns that occurred in any other cluster were counted as misses and reduced completeness. We also measured overall coverage by assessing how many words of each category were classified by the analysis.

In a similar way, we measured the extent to which information only about the preceding word predicted the syntactic category of the following word (Monaghan & Christiansen, 2004). We took the 45 most frequent *A_* frames from the Theakston et al. corpus and clustered together all words that occurred after that word (i.e., as an *x* in *Ax*). We assessed accuracy and completeness in the same way as the *AxB* analysis.

Results and Discussion

In the Theakston et al. (2001) corpus, there were a total of 1,249 different nouns and 655 verbs. For the *AxB* and *Ax* analyses, the accuracy and completeness results are shown in table 3.5. For the *AxB* analysis, a greater proportion of verbs were classified, $\chi^2 = 55.36$, $p < .001$. However, accuracy and completeness were both higher for the noun classifications, $\chi^2 = 5810.51$, and $\chi^2 = 77136.21$, respectively, both $p < .001$. Frames that formed clusters including nouns tended to contain only nouns, whereas the verb clusters tended to contain a greater proportion of words from other categories. To a greater degree, nouns tended to occur in the same clusters, whereas verbs were distributed across more clusters.

The *Ax* analysis resulted in a high coverage, classifying 97.4% of the nouns and 77.9% of the verbs—nearly six times more nouns and twice as many verbs as the *AxB* analysis (though with lower degrees of accuracy). Thus, a greater proportion of nouns than verbs were classified in the *Ax* analysis, and accuracy and completeness were both significantly higher for nouns than verbs, all $\chi^2 > 190$, $p < .001$. The clusters containing nouns contained a mean of 32.7 nouns, whereas those containing verbs contained a mean of 20.5 verbs.

Experiments 1 and 2 together indicate that the three different distributional analyses all point to an advantage for classifying nouns over verbs. The clustering of words in *AxB* and *Ax* frames is more coherent and accurate for nouns than verbs, as reflected in the advantages for accuracy and completeness analyses in Experiment 2. Nouns tend to be more frequent than verbs, in that more of the most frequent words in a corpus are nouns, in particular, proper nouns. Such a difference

Table 3.5 Accuracy and completeness of AxB and Ax frames analyses for nouns and verbs

AxB	Nouns	Verbs
Total classified	206	205
Accuracy	78%	69%
Completeness	8%	3%
Ax	Nouns	Verbs
Total classified	1217	510
Accuracy	69%	23%
Completeness	9%	3%

in frequency goes some way to explaining why distributional information may be more useful for nouns than verbs—the more information available about the context of a word, the more beneficial that information is going to be. However, frequency cannot account for differences in accuracy of classification. The distributional cues for nouns also supply more detailed information about usage than those provided for verbs. We next discuss why this points towards the greater reliance upon and benefit of word-internal cues for classifying verbs.

General Discussion

In this chapter, we have argued that language-internal information is crucial for solving the bootstrapping problem facing children in early syntax acquisition. The integration of multiple language-internal cues allows young children to get an initial grasp on lexical categories, providing a probabilistic foundation for subsequent acquisition of the syntactic constraints on their native language. Our two experiments have further illuminated the differential contribution of phonological and distributional information to the learning of nouns and verbs from English child-directed speech. Experiment 1 showed that whereas both types of cues were helpful for the discovery of nouns, phonological cues were less useful once distributional cues had been taken into account. Phonological information, on the other hand, was particularly useful for classifying verbs with distributional information providing little additional benefit to overall classification. Experiment 2 further demonstrated that these results were not due to the specific type of distributional information used in Experiment 1. Using the distributional models of Monaghan and Christiansen (2004) and Mintz (2003), we found that whereas more enriched distributional information benefits the classification of both types of words, nouns benefit considerably more than verbs both in terms of accuracy and completeness of categorization. Taken together, Experiments 1 and 2 provide

compelling evidence that phonological and distributional cues take on different, partially complementary roles in facilitating the discovery of nouns and verbs in English, with distributional information being more useful for the classification of nouns and phonological information more useful for the classification of verbs.

Our analyses in Experiment 2 further indicate that distributional information is most supportive when early verb learning is concentrated only on certain contexts initially. This is consistent with Tomasello's (1992, 2000) item-based perspective on early syntactic development, suggesting that early word learning, in particular of verbs, is centered on specific words or phrases rather than lexical categories (e.g., *eat* _____, *draw* _____, etc.). Underscoring their importance for verb learning, phonological cues may play an important role in drawing together the disparate verb fragments that were initially highlighted in an item-based fashion by distributional means. Thus, verbs need phonological cues to draw them together into a coherent lexical category because of the item-based nature of their distributional cues. Nouns, on the other hand, do not need phonological cues for this purpose because each noun tends to occur in many more different distributional contexts.

The importance of phonological cues for learning verbs becomes even more evident when language-external information is taken into account. As pointed out by Childers and Tomasello (chapter 12, this volume), nouns tend to be conceptually easier for children to learn in comparison with verbs. Understanding verbs is complicated by the fact that a particular action sequence can be conceptualized in many different ways, and determining which one is the intended one requires the child to determine what perspective a speaker is imposing on the event—the so-called packaging problem (Gentner & Boroditsky, 2001; Gleitman, 1990; Tomasello, 1992—see also chapter 8, this volume). This means that language-external cues are going to be more helpful for discovering nouns than verbs, placing even more importance on phonology as a cue to verb learning.

Not only may nouns be easier to learn because of the conceptual support of language-external cues, but they may also be easier to segment from the speech stream. The same local distributional cues that we found to be reliable for noun classification in Experiments 1 and 2 are also likely to be useful for segmenting nouns from fluent speech. Because good cues—such as *the*, *a*, and *you*—have very low transitional probabilities after them, the onsets of the nouns that tend to follow them will be easy to detect. In this way, the same distributional properties that make nouns easy to classify will also make them easier to discover in the speech stream: nouns are more predictable because they follow good distributional cues. This would suggest a disadvantage for segmenting verbs in early acquisition, and may thus explain the developmental delay in verb segmentation relative to noun segmentation observed by Nazzi and Houston (chapter 2, this volume).

Although we have uncovered several factors that are likely to make verb learning in English more difficult than the learning of nouns, it may be that it is

not verbs per se that are hard to learn, but words that are semantically complex. There is an ongoing debate concerning whether nouns are harder to learn than verbs (see chapters 12, 17, and 18 for discussions). A possible consensus is that it is the concreteness of a word that determines its ease of learning, rather than its lexical category (e.g., chapter 14 this volume). Our approach to multiple-cue integration accommodates this perspective. Concrete words will have more reliable language-external cues than more abstract words, and this will facilitate word learning through multiple-cue integration. Words that have a more reliable set of cues supporting them are going to be easier to learn—no matter whether the cues are language-internal or language-external. However, when it comes to associating actions with words, distributional and semantic cues are generally going to be less useful, leaving phonological cues as the perhaps most important cue for the discovery of verbs through multiple-cue integration.

Acknowledgments This research was supported in part by a Human Frontiers Science Program Grant (RGP0177/2001-B). Thanks to Rick Dale, Thomas Farmer, Stanka Fitneva, Roberta Golinkoff, and Luca Onnis for their comments on an earlier version of this chapter.

References

- Baayen, R. H., Pipenbrock, R., & Gulikers, L. (1995). *The CELEX Lexical Database* (CD-ROM). Philadelphia: Linguistic Data Consortium, University of Pennsylvania.
- Bijeljic, R., Bertocini, J., & Mehler, J. (1993). How do 4-day-old infants categorize multisyllabic utterances? *Developmental Psychology*, 29, 711–721.
- Billman, D. (1989). Systems of correlations in rule and category learning: Use of structured input in learning syntactic categories. *Language and Cognitive Processes*, 4, 127–155.
- Brooks, P. J., Braine, M. D., Catalano, L., & Brody, R. E. (1993). Acquisition of gender-like noun subclasses in an artificial language: The contribution of phonological markers to learning. *Journal of Memory and Language*, 32, 76–95.
- Campbell, R., & Besner, D. (1981). This and thap—Constraints on the pronunciation of new written words. *Quarterly Journal of Experimental Psychology*, 33, 375–396.
- Cassidy, K. W., & Kelly, M. H. (1991). Phonological information for grammatical category assignments. *Journal of Memory and Language*, 30, 348–369.
- Cassidy, K. W., & Kelly, M. H. (2001). Children's use of phonology to infer grammatical class in vocabulary learning. *Psychonomic Bulletin and Review*, 8, 519–523.
- Christiansen, M. H., & Dale, R. (2001). Integrating distributional, prosodic and phonological information in a connectionist model of language acquisition. In *Proceedings of the 23rd annual conference of the Cognitive Science Society* (pp. 220–225). Mahwah, NJ: Lawrence Erlbaum.
- Cutler, A. (1993). Phonological cues to open- and closed-class words in the processing of spoken sentences. *Journal of Psycholinguistic Research*, 22, 109–131.
- Cutler, A., & Carter, D. M. (1987). The predominance of strong initial syllables in the English vocabulary. *Computer Speech and Language*, 2, 133–142.
- Dunning, T. (1993). Accurate methods for the statistics of surprise and coincidence. *Computational Linguistics*, 19, 61–74.
- Fernald, A., & McRoberts, G. (1996). Prosodic bootstrapping: A critical analysis of the argument and the evidence. In J. L. Morgan & K. Demuth (Eds.), *Signal to syntax: Bootstrapping from speech to grammar in early acquisition* (pp. 365–388). Mahwah, NJ: Lawrence Erlbaum.
- Fisher, C., & Tokura, H. (1996). Acoustic cues to grammatical structure in infant-directed speech: Cross-linguistic evidence. *Child Development*, 67, 3192–3218.
- Frigo, L., & McDonald, J. L. (1998). Properties of phonological markers that affect the acquisition of gender-like subclasses. *Journal of Memory and Language*, 39, 218–245.
- Gentner, D., & Boroditsky, L. (2001). Individuation, relational relativity and early word learning. In M. Bowerman & S. Levinson (Eds.), *Language acquisition and conceptual development* (pp. 215–256). Cambridge: Cambridge University Press.
- Gerken, L. A. (1996). Prosody's role in language acquisition and adult parsing. *Journal of Psycholinguistic Research*, 25, 345–356.
- Gerken, L. A., Jusczyk, P. W., & Mandel, D. R. (1994). When prosody fails to cue syntactic structure: Nine-month-olds' sensitivity to phonological vs. syntactic phrases. *Cognition*, 51, 237–265.
- Gleitman, L. (1990). The structural sources of verb meanings. *Language Acquisition*, 1, 3–55.
- Gleitman, L., & Wanner, E. (1982). Language acquisition: The state of the state of the art. In E. Wanner & L. Gleitman (Eds.), *Language acquisition: The state of the art* (pp. 3–48). Cambridge: Cambridge University Press.
- Gómez, R. L., & Gerken, L. A. (1999). Artificial grammar learning by 1-year-olds leads to specific and abstract knowledge. *Cognition*, 70, 109–135.
- Gómez, R. L., & Gerken, L. A. (2000). Infant artificial language learning and language acquisition. *Trends in Cognitive Sciences*, 4, 178–186.
- Green, T. R. G. (1979). The necessity of syntax markers: Two experiments with artificial languages. *Journal of Verbal Learning and Verbal Behavior*, 18, 481–496.
- Jusczyk, P. W. (1997). *The discovery of spoken language*. Cambridge, MA: MIT Press.
- Jusczyk, P. W. (1999). How infants begin to extract words from speech. *Trends in Cognitive Sciences*, 3, 323–328.
- Jusczyk, P. W., & Kemler-Nelson, D. G. (1996). Syntactic units, prosody, and psychological reality during infancy. In J. L. Morgan & K. Demuth (Eds.), *Signal to syntax: Bootstrapping from speech to grammar in early acquisition* (pp. 389–408). Mahwah, NJ: Lawrence Erlbaum.
- Karmiloff-Smith, A. (1979). *A functional approach to child language: A study of determiners and reference*. Cambridge: Cambridge University Press.
- Kelly, M. H. (1988). Phonological biases in grammatical category shifts. *Journal of Memory and Language*, 27, 343–358.
- Kelly, M. H. (1992). Using sound to solve syntactic problems: The role of phonology in grammatical category assignments. *Psychological Review*, 99, 349–364.
- Kelly, M. H., & Bock, J. K. (1988). Stress in time. *Journal of Experimental Psychology: Human Perception and Performance*, 14, 389–403.
- Kemler Nelson, D. G., Hirsh-Pasek, K., Jusczyk, P. W., & Wright Cassidy, K. (1989). How the prosodic cues in motherese might assist language learning. *Journal of Child Language*, 16, 55–68.
- Kuhl, P. K. (1999). Speech, language, and the brain: Innate preparation for learning. In M. D. Hauser & M. Konishi (Eds.), *The design of animal communication* (pp. 419–450). Cambridge, MA: MIT Press.

- Kuhl, P. K., Andruski, J. E., Chistovich, I. A., Chistovich, L. A., Kozhevnikova, E. V., Ryskina, V. L., et al. (1997). Cross-language analysis of phonetic units in language addressed to infants. *Science*, 277, 684–686.
- MacWhinney, B. (2000). *The CHILDES Project: Tools for analyzing talk* (3rd ed.). Mahwah, NJ: Lawrence Erlbaum.
- Maratsos, M., & Chalkley, M. A. (1980). The internal language of children's syntax: The ontogenesis and representation of syntactic categories. In K. Nelson (Ed.), *Children's language* (Vol. 2, pp. 127–214). New York: Gardner Press.
- Marchand, H. (1969). *The categories and types of present-day English word-formation* (2nd ed.). Munich: C.H. Beck'sche Verlagsbuchhandlung.
- Mattys, S. L., Jusczyk, P. W., Luce, P., & Morgan, J. L. (1999). Phonotactic and prosodic effects on word segmentation in infants. *Cognitive Psychology*, 38, 465–494.
- McDonald, J. L., & Plauche, M. (1995). Single and correlated cues in an artificial language learning paradigm. *Language and Speech*, 38, 223–236.
- Mehler, J., Jusczyk, P. W., Lambert, G., Halsted, N., Bertoncini, J., & Amiel-Tison, C. (1988). A precursor of language acquisition in young infants. *Cognition*, 29, 143–178.
- Meier, R. P., & Bower, G. H. (1986). Semantic reference and phrasal grouping in the acquisition of a miniature phrase structure language. *Journal of Memory and Language*, 25, 492–505.
- Mintz, T. H. (2003). Frequent frames as a cue for grammatical categories in child directed speech. *Cognition*, 90, 91–117.
- Mintz, T. H., Newport, E. L., & Bever, T. G. (2002). The distributional structure of grammatical categories in speech to young children. *Cognitive Science*, 26, 393–424.
- Monaghan, P., Chater, N., & Christiansen, M. H. (2005). *The differential contribution of phonological and distributional cues in grammatical categorization*. Manuscript submitted for publication.
- Monaghan, P., & Christiansen, M. H. (2004). What distributional information is useful and usable for language acquisition? In *Proceedings of the 26th annual conference of the Cognitive Science Society* (pp. 963–968). Mahwah, NJ: Lawrence Erlbaum.
- Morgan, J. L. (1996). Prosody and the roots of parsing. *Language and Cognitive Processes*, 11, 69–106.
- Morgan, J. L., & Demuth, K. (Eds.). (1996). *Signal to syntax: Bootstrapping from speech to grammar in early acquisition*. Mahwah, NJ: Lawrence Erlbaum.
- Morgan, J. L., Meier, R. P., & Newport, E. L. (1987). Structural packaging in the input to language learning: Contributions of prosodic and morphological marking of phrases to the acquisition of language. *Cognitive Psychology*, 19, 498–550.
- Morgan, J. L., & Saffran, J. R. (1995). Emerging integration of sequential and suprasegmental information in preverbal speech segmentation. *Child Development*, 66, 911–936.
- Morgan, J. L., Shi, R., & Allopenna, P. (1996). Perceptual bases of grammatical categories. In J. L. Morgan & K. Demuth (Eds.), *Signal to syntax: Bootstrapping from speech to grammar in early acquisition* (pp. 263–283). Mahwah, NJ: Lawrence Erlbaum.
- Newport, E. L., Gleitman, H., & Gleitman, L. R. (1977). Mother, I'd rather do it myself: Some effects and non-effects of maternal speech style. In C. E. Snow & C. A. Ferguson (Eds.), *Talking to children: Language input and acquisition* (pp. 109–149). Cambridge: Cambridge University Press.
- Pallier, C., Christophe, A., & Mehler, J. (1997). Language-specific listening. *Trends in Cognitive Sciences*, 1, 129–132.
- Pinker, S. (1984). *Language learnability and language development*. Cambridge, MA: Harvard University Press.

- Real, F., Christiansen, M. H., & Monaghan, P. (2003). Phonological and distributional cues in syntax acquisition: Scaling up the connectionist approach to multiple-cue integration. In *Proceedings of the 25th annual conference of the Cognitive Science Society* (pp. 970–975). Mahwah, NJ: Lawrence Erlbaum.
- Redington, M., Chater, N., & Finch, S. (1998). Distributional information: A powerful cue for acquiring syntactic categories. *Cognitive Science*, 22, 425–469.
- Redington, M., Chater, N., Huang, C. R., Chang, L. P., Finch, S., & Chen, K. J. (1995). The universality of simple distributional methods: Identifying syntactic categories in Mandarin Chinese. In *Proceedings of the Cognitive Science of Natural Language Processing*. Dublin: Dublin City University.
- Saffran, J. R. (2003). Statistical language learning: Mechanisms and constraints. *Current Directions in Psychological Science*, 12, 110–114.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926–1928.
- Sereno, J. A., & Jongman, A. (1990). Phonological and form class relations in the lexicon. *Journal of Psycholinguistic Research*, 19, 387–404.
- Sereno, J. A., & Jongman, A. (1995). Acoustic correlates of grammatical class. *Language and Speech*, 38, 57–76.
- Shady, M., & Gerken, L. A. (1999). Grammatical and caregiver cues in early sentence comprehension. *Journal of Child Language*, 26, 163–175.
- Shafer, V. L., Shucard, D. W., Shucard, J. L., & Gerken, L. A. (1998). An electrophysiological study of infants' sensitivity to the sound patterns of English speech. *Journal of Speech, Language, and Hearing Research*, 41, 874–886.
- Shi, R., Morgan, J., & Allopenna, P. (1998). Phonological and acoustic bases for earliest grammatical category assignment: A cross-linguistic perspective. *Journal of Child Language*, 25, 169–201.
- Shi, R., Werker, J. F., & Morgan, J. L. (1999). Newborn infants' sensitivity to perceptual cues to lexical and grammatical words. *Cognition*, 72, B11–B21.
- Steinhauer, K., Alter, K., & Friederici, A. D. (1999). Brain potentials indicate immediate use of prosodic cues in natural speech processing. *Nature Neuroscience*, 2, 191–196.
- Theakston, A. L., Lieven, E. V. M., Pine, J. M., & Rowland, C. F. (2001). The role of performance limitations in the acquisition of verb-argument structure: An alternative account. *Journal of Child Language*, 28, 127–152.
- Tomasello, M. (1992). *First verbs: A case study of early grammatical development*. New York: Cambridge University Press.
- Tomasello, M. (2000). The item-based nature of children's early syntactic development. *Trends in Cognitive Sciences*, 4, 156–163.
- Tomasello, M., Kruger, A., & Ratner, H. (1993). Cultural learning. *Behavioral and Brain Sciences*, 16, 495–552.
- Valian, V., & Coulson, S. (1988). Anchor points in language learning: The role of marker frequency. *Journal of Memory and Language*, 27, 71–86.
- Valian, V., & Levitt, A. (1996). Prosody and adults' learning of syntactic structure. *Journal of Memory and Language*, 35, 497–516.
- Weissenborn, J., & Höhle, B. (Eds.). (2001). *Approaches to bootstrapping: Phonological, lexical, syntactic and neurophysiological aspects of early language acquisition*. Philadelphia: John Benjamins.
- Werker, J. F., & Tees, R. C. (1999). Influences on infant speech processing: Toward a new synthesis. *Annual Review of Psychology*, 50, 509–535.

Action Meets Word

How Children Learn Verbs

EDITED BY
Kathy Hirsh-Pasek
Roberta Michnick Golinkoff

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Oxford University Press, Inc., publishes works that further
Oxford University's objective of excellence
in research, scholarship, and education.

Oxford New York
Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in
Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Copyright © 2006 by Oxford University Press, Inc.

Published by Oxford University Press, Inc.
198 Madison Avenue, New York, New York 10016

www.oup.com

Oxford is a registered trademark of Oxford University Press

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
electronic, mechanical, photocopying, recording, or otherwise,
without the prior permission of Oxford University Press.

Library of Congress Cataloging-in-Publication Data
Action meets word : how children learn verbs / edited by Kathy Hirsh-Pasek
and Roberta Michnick Golinkoff.

p. cm.

Includes bibliographical references and index.

ISBN 0-19-517000-8

ISBN-13 978-0-19-517000-9

1. Language acquisition. 2. Grammar, Comparative and general--Verb.

I. Hirsh-Pasek, Kathy. II. Golinkoff, Roberta M.

P118.A174 2005

401'.93--dc22 2005008496

9 8 7 6 5 4 3 2 1