In A. Greenhill, M. Hughs, H. Littlefield, & H. Walsh (eds.), *Proceedings of the Twenty-third Boston University Conference on Language Development*. Somerville, MA: Cascadilla Press (1999).

The Successes and Failures of Word-to-World Mapping

Jesse Snedeker and Lila Gleitman
University of Pennsylvania/IRCS

Michael Brent
The Johns Hopkins University

Introduction

This research examines how children begin learning the meanings of words. Necessarily, novice word-learners must start by pairing a word form with the scenes in which it occurs, collecting several such pairs, and then identifying the common element in the scenes. Before children have learned many words or any language specific syntax, this extralinguistic context is the only information source they can draw on. However, there has been relatively little research on the details of cross-situational observation. We do not know just how much information is actually available in the extralinguistic contexts of word use. Nor do we know how the properties of this information source might shape what children learn in the initial stages of vocabulary acquisition or the process by which they learn it. The studies that follow begin to address these questions. They focus exclusively on the informational content of the input. They do not explore whether this information is accessible to infants.

If all words were equally easy to learn, we'd expect that verbs would be learned before nouns. Both adult-to-adult and infant directed speech contains fewer tokens of each type of noun and more tokens of each type of verb (Gentner & Boroditsky, in press). If learning depends on the number of exposures to the material to be learned, then children should first learn the words they hear most frequently. However, in early vocabularies nouns predominate while verbs are under-represented. This is true not only in languages like English (Brown 1973; Bates, Dale & Thal 1995) in which subjects and objects are mandatory and verbs are often buried in the middle of an utterance, but also in languages where verbs may be more salient such as: Italian (Casselli et al., 1995); Chinese (Tardif & Gelman, 1998); Korean (Au, Dapretto & Song, 1994) and many others (Gentner, 1982).

This pattern is often attributed to the conceptual limitations of young children (Gentner, 1978; Macnamara, 1972). The argument goes something like this: the novice word learner has only simple concepts (such as those specifying

individual persons and categories of perceptually similar objects) and therefore can only learn the words that encode these simple concepts, which are mostly nouns. As the child develops, she becomes capable of entertaining more complex concepts (such as relations, events, or actions) and thus armed can begin to map these concepts to the words that express them (which are mostly verbs and prepositions). The view that conceptual growth sets the pace for lexical learning has been so influential that many researchers have used changes in vocabulary composition as measures of conceptual change (see e.g., Smiley & Huttenlocher, 1995).

Gillette, Gleitman, Gleitman & Lederer (in press, henceforth GGGL) have offered an alternate explanation. According to these authors, the selective learning of nouns is due to the properties of the input. Because the initial information source for word learning, extralinguistic context, provides far more information about the meanings of nouns than the meanings of verbs, children learn mostly nouns until other sources of information become available.

Human Simulations

GGGL's claim is about the relationship between two types of information: the properties of the scene and aspects of word meaning. Computer simulations are useful for studying questions of this type. For example, Brent (1994) used a computer simulation to examine whether the presence of particular lexical items correlates with the syntactic frames in which a verb appears in child-directed speech (see also Fisher, Gleitman & Gleitman, 1991 for an experimental demonstration). Simulations of this kind are not designed to mimic the mechanisms by which infants learn language. They simply perform a correlational analysis to test whether one type of information is a good predictor of another type. Any machine capable of representing both types of information can be used to conduct this kind of analysis. In these studies, we will be using adult human beings. Adults can represent both scenes and word meanings. While their representations may differ from those of small children, there is no reason to believe that these differences are more dramatic than the differences between an infant and the coding system that underlies the representations of any computer simulation. Infants, after all, grow up to be adults.

GGGL presented the word-scene pairs from infant-directed speech to college students. The target words were common nouns and verbs drawn from transcripts of videotapes of mothers playing with their children (18-24 months). Subjects were told the syntactic category of each of the words and were shown six short video clips of scenes in which the word was used. The audio had been removed from the video clips and the word of interest was replaced with a beep. The subjects' task was to guess the word that the mother was saying. This study provided a test of the conceptual limitation hypothesis by presenting the actual input to fully mature learners. College students, like toddlers, were able to learn more nouns than verbs. On the final trial, they identified 45% of the noun targets but only 15% of the verbs. The failure of adults to learn common verbs

such as *think, make,* and *play* cannot be attributed to their inability to entertain the concepts these words encode. Presumably, they failed because the scenes did not provide enough information to disambiguate the meaning of the beep.

The purpose of the GGGL was to determine how much a novice could learn solely on the basis of cross-situational observation. To test this, they attempted to remove all other cues to word meaning; for example, by removing the audio they eliminated information about the words that co-occurred with the target or the structure of that sentence. However, they gave their subjects one additional piece of information: the subjects were told the syntactic category of the target word. This is information which novice language learners certainly do not have. Syntactic categories are not directly available in the input, and to discover the language specific cues of syntactic class a child must, at minimum, first learn the meanings of some words from the relevant classes so that he can correctly label the category he has identified.

GGGL tests the potency of a combination of two information sources: word-to-world pairs and syntactic category labels. It is not clear, a priori, how this second information type would interact with the first. In the study that follows we remove this additional information, to place the adult subjects in a position more like that of actual novice word learners. We hope to answer the following questions: 1) Would subjects be able to correctly identify words if they didn't know their category? and 2) Would their performance on nouns and verbs be affected similarly—was the noun advantage in GGGL attributable to an interaction between the two information types?

Experiment 1

To answer these questions we made two changes to the protocol used in GGGL. First, subjects were given no information about the syntactic category of a given target. Second, they were told that half the words were nouns and half were verbs. This was done to ensure that subjects expectations of the type of words that they might encounter would be accurate.

Subjects were tested in groups of one to four. They were told that they would be seeing short silent videos of mothers playing with their young children and that their task was to identify a word that the mother had said during that video. The unknown word was represented by a beep that occurred at the exact point at which the mother had used the word. For each word subjects saw a set of scenes that included six instances of that word. After each beep they were asked to guess the identity of the word based on the scenes that they had seen thus far. Finally, after they made their sixth guess, the tape was paused for 10 seconds and subjects were encouraged to look over all of their guesses and make a final effort to identify the word. This procedure was repeated for 16 different words, 8 nouns and 8 verbs. Before the study began, subjects were told that the children were 18-24 months of age, that the words had been chosen for their frequency in speech to young children, that roughly half of them were nouns and half verbs, and that none were exclamations or social words such as *hello* or

thank you. The 84 undergraduate subjects were randomly assigned to one of three stimulus lists and to one of two presentation orders. They received either partial class credit in a psychology course or a small payment for their participation.

The videotaped stimuli were taken from GGGL (a more complete description of their construction is provided therein). The stimuli were drawn from 4 hour-long videotaped play sessions. In each session, a mother and her 18-24 month old child explored a bag of toys that the experimenter had brought to the child's home. The 24 most common nouns and 24 most common verbs from these tapes were chosen as targets. The target words were divided between three lists, each containing an equal number of nouns and verbs. The order of presentation of the targets in the lists was randomly generated and this order of presentation was reversed for half of the subjects who received each list.

For each target word, six instances in which the word was used were selected for inclusion in the study. Uses of a target word were excluded if the referent of the word was visible to the child but off camera, or if the mothers lips were visible and might provide information about the word form. When there were more than six instances of the word that met these criteria, instances were selected at random. For each instance, a video clip was constructed that began 30 seconds before the target word was used and ended 10 seconds after it was said. In many cases, the mother said the target word at another time during this 40-second period. In these cases, the clip was expanded to include 30 seconds before the first use and 10 seconds after the last. Each of the uses of the word in a single clip counted as one of the six stimuli.

Results

Responses were scored as correct only if they contained the same base morpheme as the target word. Guesses that differed from the target in number, tense, or voice were coded as correct. The number of correct responses to the final trial was analyzed both by subjects and by items. In both analyses, there was a significant effect of syntactic category ($p < .001$, $p < .05$). Subjects correctly identified nouns on 26% of the final trials but correctly identified verbs on only 12%. The subjects attempted to identify each target word seven times. There was a significant effect of trial on number of correct responses, signifying that subjects were more likely to learn a word after they had seen more scenes in which it was used ($p < .001$, $p < .001$). There was an interaction between trial and syntactic category in the subjects analysis, indicating that this improvement was more pronounced for nouns than for verbs ($p < .05$).

While subjects identified more nouns than verbs, there was also substantial variation in performance on items within a syntactic class. For example, 79% of the subjects were able to identify the noun *plane* but no one identified *things*. What accounts for this variability? Our subjects appeared to succeed on precisely those words that young children are also able to learn: concrete or imageable words, especially those denoting basic-level object categories. Item

performance correlates with ratings of imageability (collected by GGGL). Basic-level object terms were identified more than four times as often as other nouns. Verbs that denote mental or perceptual states or acts of communication were identified less than half as often as verbs denoting actions.

The percentage of correct guesses tells us whether the scene provides cues to the target word but does not tell us whether the scene is also providing misleading cues that point towards other conjectures. Does extralinguistic context ever lead subjects astray? Is misleading information more of a problem for nouns or for verbs? These questions were addressed by: 1) identifying the strongest competitor for each of the target words (this was defined as the non-target word that was selected by the greatest number of subjects) and 2) comparing the number of subjects who chose the target versus the number who chose the competitor for both nouns and verbs. There was a significant interaction between type of response and syntactic category of target ($p < .01$). Targets were favored for nouns. But for verbs, subjects were more likely to choose the strongest competitor than the target.

These results confirm and strengthen the conclusions of GGGL. Even without the supplementation of syntactic category labels, word-to-world pairings provide reasonably good information about the meanings of nouns but they provide misleading information about the meanings of verbs. The pattern of early vocabulary acquisition can be accounted for by the properties of the input. It is therefore unnecessary to posit a conceptual deficit in the learner. Children learn nouns because word-to-world pairs unveil the meanings of nouns. They cannot learn verbs until they can make use of other types of information, such as noun co-occurrence and the syntactic frame in which the verb appears (for which, see GGGL).

Are syntactic category labels a potentially useful information source? Did the additional information provided in GGGL help the subjects to learn targets of a particular type? To address these questions, we analyzed Experiment 1 and GGGL as two conditions in a single experiment. With the exception of the changes described above, the two studies were identical; both studies used the same stimuli, instructions, and subject pool. The numbers of correct responses to the final guess for both experiments were entered into a subject and an item ANOVA. There was a significant effect of experiment ($p < .001$, $p < .001$); the presence of the category label in GGGL helped subjects to identify the target word. But this was superseded by an experiment by syntactic category interaction ($p < .001$, $p < .001$); category labels substantially improved performance on nouns (from 25 to 45%) but had relatively little effect on verbs (from 12 to 15%)

Capitalizing on the Noun Advantage

For our subjects, knowing that they were looking for a noun helped them identify the word, while knowing that they were looking for a verb did not. This suggests an adaptive strategy for learners confronted with data of this type:

Adopt a bias to look for meanings that could potentially be carried by nouns. By doing this they will increase their chances of correctly discovering the meaning of the word. When subjects made a guess that was a noun, they were right 37% of the time. When they made a verb guess, they were right only 10% of the time. But what does it mean to look for a meaning that could be a noun? Among other things, nouns can pick out classes of objects, persons, events, actions and properties. Attempting to focus one's search by looking for all of these things at once doesn't seem very efficient. But although the stimuli in Experiment 1 were identified and divided by syntactic category, there is no reason to insist that this distinction is the best explanation for the differences in learnability between the two classes of words. Many have argued that specific syntactic categories and semantic categories are consistently coupled across languages. The best (and perhaps the only truly universal) example of such a linkage is the inclusion of words denoting persons or classes of objects in the category of nouns (Maratsos, 1991). Seventy-five percent of the nouns in Experiment 1 were labels for classes objects or people. Subjects did twice as well on these words as they did on the nouns that denoted abstractions, actions, or parts. Thus, the best strategy for learner relying on word-to-world pairs might be to initially attempt to map every new word to a class of objects or persons.

This idea, of course, is nothing new. Proposals and evidence for closely related strategies abound (Macnamara, 1982; Markman & Wachtel 1988; Golinkoff, Mervis, & Hirsh-Pasek, 1994). This suggestion is probably most similar to Markman's Whole-Object Constraint, the proposal that children initially "assume that a novel label is likely to refer to the whole object and not to its parts, substance, or other properties" (Markman, 1992, p.61). However, defined in this way the Whole Object Constraint doesn't address the issue of how a child discovers that a word is linked to a particular object rather than to a relation between objects or an action (typically denoted by a verb). This is in large part a reflection of the type of studies that motivated the constraint. While the human simulations provide authentic input to unauthentic "infants", word-learning studies of this type provide simplified "input" to authentic children. An adult points to a stationary object and utters a nonsense word several times. Next, new objects are brought out which differ from the standard in varying ways and child is asked—indirectly of course—whether the new label can apply to these objects as well. The child's response is constrained both by the situation in which the word is originally used and by the set of alternatives that are presented. Since no salient action or relation is offered or tested, these studies cannot determine whether children prefer to map words to "nouny" referents over "verby" referents.

Is there any evidence that children have such a preference? The composition of early vocabulary provides some hints. A noun bias is certainly consistent with the disproportionate representation of nouns in early vocabularies. However, if there is such a bias, it appears to build as the child learns more words. The first 10 or 20 words that a child learns typically come

from a broad range of categories (Benedict, 1979). While nouns are still more common than words of any other category, they often make up less than half the words at this stage and variation across children is substantial (Nelson, 1973). The proportion of nouns increases as children learn more words, until it levels off around the time the child has a vocabulary of 200 words (Bates, Dale & Thal, 1995).

Experimental evidence for a noun bias is also quite limited. Woodward (1992) found that by 18 months children selectively attend to objects rather than moving substances in response to a novel label. This suggests that by this age learners would rather map a word to an object than to an action. However, infant habituation studies suggest that this bias only develops at around 13 months, approximately the age at which children begin to produce their first words (Echols, 1990). Prior to this, children appear to link a novel label to that aspect of a situation (action or object) which occurs most reliably in its presence. Studies of controlled, naturalistic, word learning have found that, while all children learn object labels with fewer exposures than action labels, this difference is heightened in children who have already acquired a fair number of nouns (Leonard, Schwartz, Morris & Chapman, 1981). In sum, experimental work on the emergence of a noun bias is still somewhat inconclusive, but suggests that the bias is not present at the onset of vocabulary acquisition. Rather, it appears soon after.

Where does this noun bias come from? Three possible origins for word-learning constraints have been suggested in the literature. First, the bias could be an innate constraint that matures after word learning is underway or is triggered by limited exposure to language. Second, the constraint might be based on knowledge from other domains (perception, attention, and categorization). When adequate learning has occurred in these realms then the constraint becomes available to guide word learning. This is what we take to be Markman's position with regards to the whole object constraint (1992). Finally, it is possible that the constraint develops as a result of word learning, that it emerges from an interaction between the child and the input. This position has been developed and advanced by L. Smith in work on the origins of the shape bias (1995).

It is this third position that we explore by offering a scenario for how the noun bias could be learned in the input situation of early word learning. Children begin lexical acquisition without a noun bias: They gamely entertain the open-ended hypothesis that a word can stand for any of the things that words often do. The word learning system, however, can monitor its success—either because it receives explicit feedback in the form of corrections or failures to communicate or because it gets implicit feedback from the further situations in which the word occurs. Now as we saw earlier, because of the properties of word-to-world pairs, novices will tend to succeed when they make a noun guess but fail when they make a verb guess. The word-learning system exploits this by adopting a bias in favor of nominal meanings (people and objects). The comparison of word-to-world mapping with and without category labels suggests that this

strategy will lead to an even more efficient learning of nouns and an increasingly nominal vocabulary. Verbs, adjectives, and prepositions have to take a back seat until other information sources became available.

In our story, the shift to a noun bias is brought about by experience with a particular type of input to word learning. In the other two tales offered above, the plot was driven by the almost inevitable changes brought on by maturation or general cognitive development. The alternatives are in many ways parallel to the explanations that were offered for the over-representation of nouns in early vocabularies. In both cases, accounts relying on the cognitive status of the learner are pitted against accounts that depend on specific properties of the input. The human simulation paradigm allows us to hold the input constant and change the maturational status of the learner, thus disentangling the two possible causes. In the next experiment we used this technique to examine: 1) whether cognitively mature subjects would shift towards a noun bias as they gained experience with the input; and 2) if so, what kind of evidence about performance would be necessary to produce this shift: Is implicit information available or do subjects require explicit feedback?

Experiment 2

Thirty-two subjects were randomly assigned to one of two conditions (crossed by two presentation orders) and tested individually. In the No-Feedback condition, the procedure was essentially the same as that described in Experiment 1: subjects saw six instances of each target word, made a guess after each instance, then made a final guess, and moved on to the next target. The subjects in the Feedback condition were told after their final guess whether they had correctly identified the word. All of the subjects saw videos for the same 24 target words. Subjects were told that the target words were nouns and verbs but since we were interested in the proportion of noun guesses, we did not constrain them by telling them the relative numbers of each. To increase subject motivation we rewarded every correct final guess with 50¢. Subjects were told about the reward before the study began but were not told how much they had earned until the study was completed. This reward was in addition to the compensation described above.

The target words were the 12 most common nouns and the 12 most common verbs from Experiment 1, and the video clips were the same ones that were used in that study. Because the questions of interest are about changes in strategy as the study unfolds, it was critical to balance the difficulty and syntactic category of the words across the list. This was done by dividing both the nouns and verbs into three groups based on performance on these words in Experiment 1. One noun and one verb from each group were randomly assigned to one quarter of the new stimuli list. The position of the words in the first quarter of the list was generated randomly and this segment was used as a template for the other three. When these four segments were joined, they made up the first order. The second order was made by reversing the position of the

segments. Thus the words that made up the first quarter for half the subjects were the words that made up the last quarter for the other half and occurred in the same order for both groups, guaranteeing that comparisons between the first and final quarters will not be polluted by item or order effects.

After the subject had completed all 24 words, the experimenter went through the subject's responses looking for words that could potentially belong to more than one syntactic category. Subjects were asked to decide whether the word was a noun, a verb, or a member of another category. If they were having trouble with this distinction, they were told that nouns were often persons, places, or things and that verbs were often actions, relations, or properties. The category of the word was coded based on these responses.

The dependent measure in the analyses that follow is the number of noun guesses that the subjects made for a given group of items. To answer our questions, two analyses were necessary. First, to determine whether experience with the input leads to a noun bias, we had to compare the proportion of noun guesses in the initial trials with the proportion in the final trials. Second, to test whether explicit feedback about performance is necessary for this shift in bias, we had to compare the number of noun guesses in the final trials of Feedback and No-Feedback conditions. Both a subject and an item ANOVA of the number of noun guesses to the first six trials and final six trials were performed. Experience with the task led to a significant increase in the number of noun guesses ($p < .001$, $p < .01$). Feedback, however, had no significant effect ($F < 1$ and $p > .5$ for both) and there was no interaction between the Feedback and experience with the task ($F < 1$ and $p > .3$ for both). Subjects learned to make more noun guesses as the experiment progressed regardless of whether they were given explicit information about the success of their attempts.

Experiment 3

There are two potential interpretations of this pattern of results. First, as suggested above, it is possible that subjects have access to implicit evidence about their performance and evaluate whether their responses are hitting the mark even if the experimenter doesn't help them. They might notice whether the guess they make for a given word on one trial predicts aspects of the scenes that they see for that same word on subsequent trials. Implicit information of this type might allow them to notice that noun guesses are more successful than verb guesses and push them to develop a noun bias. Second, these results might be due to an experimental artifact; exposure to the tapes or the task or the activity of thinking about words might lead subjects to make more noun guesses regardless of the success of this strategy. Distinguishing between these two possibilities requires disentangling the accumulation of implicit evidence from the amount of exposure to the task and materials. We did this by creating stimuli in which there is no correlation between the placement of the six beeps and any aspect of the scenes, thus no systematic information about word meaning. If the correlational structure of the information source is causing the

subjects to shift towards nouns, subjects exposed to these tapes should not make this shift. If, however, the bias is due to the task or familiarity with the tapes, then it should occur in this study as well.

The procedure in this study was identical to that of the No-Feedback condition in Experiment 2. Since the subjects in that study were told that they would receive 50¢ for each correct word, the subjects in this study were told the same. However, since there were no actual target words in these stimuli, the experimenter simply marked 4 random responses as correct at the end of the session and gave the subject $2. The 22 subjects were randomly assigned to one of two orders of presentation. The stimuli consisted of 24 false target words. The videotapes from Experiment 2 served as templates for these videotapes. Each false target contained as many scenes as the corresponding word in the template and the beeps were divided amongst those scenes in the same way. Each scene however, was chosen at random from the full set used in Experiment 1. The beeps were placed randomly in the scenes with the caveat that a beep not occur within 10 seconds of where one had been in the original stimuli. Thus, the video for each false target consisted of set of scenes that had originally be chosen for several different words with beeps placed at points where these words had not occurred.

The responses were coded by category as described above, for each subject the number of noun guesses for the first and the final six false targets were separately summed, and these totals were analyzed in a subject ANOVA. Experience with the task had no significant effect on the number of noun guesses ($F < 1$, $p > .5$). When the noun totals for this study were compared directly with the totals from Experiment 2, there was a significant effect of experiment; subjects in Experiment 3 made far fewer noun guesses ($p < .001$). There was also a significant interaction between experiment and level of experience ($p < .01$). While subjects who were receiving real word-scene pairs made more noun guesses as the experiment progressed, subjects given tapes that contained no information about word meaning actually made fewer. Exposure alone does not lead to a noun shift. This shift must be attributable to the correlational structure of input. Subjects make more noun guesses only when the input teaches them that this is a successful strategy.

Tentative Conclusions

This line of research is far from complete. We are claiming that people rely on implicit evidence about their performance to shape their word-learning strategies, but we have provided only indirect evidence that they can evaluate their own success at this task. We argue that children might learn from the input to prefer mapping words to objects, yet we admit that the evidence that word learning precedes the object bias is fairly weak. Clearly, we have alot of work left to do. Nevertheless, it is possible to offer some tentative conclusions based on what we have done so far. First, the predominance of nouns in early vocabularies need not be attributed to the cognitive limitations of young

children. The input (word-scene pairs) provides more evidence for the meanings of nouns than it does for the meanings of verbs. Second, the development of word-learning biases need not be attributed to spontaneous changes in the learner. Experience with word-scene pairs can shape the strategies of learners.

Notes

*We thank Jamie Schuller, Bjoern Hartmann, Le'eat Sharoni, Deborah Small, Jelena Vojinovic, Michel Walker, Simon Cheung, Karen Ram, and Daniele Spira for their assistance and Jane Gillette, Henry Gleitman, and the members of the Cheese Seminar for their advice.

References

Au, Terry; Dapretto, Mirella & Song, You-Kyung (1994). Input versus constraints: Early word acquisition in Korean and English. *Journal of Memory and Language,* **33**, 567-582.

Bates, Elizabeth; Dale, Philip & Thal, Donna (1995). Individual differences and their implication for theories of language development. In P. Fletcher & B. MacWhinney (Eds.), *Handbook of Child Language* (96-151). Oxford: Basil Blackwell.

Benedict, Helen (1979) Early lexical development: Comprehension and production. *Journal of Child Language,* **6**, 183-200.

Brent, Michael (1994). Surface cues and robust inference as a basis for the early acquisition of subcategorization frames, in L. Gleitman and B. Landau (Eds.) pp. 433-470, *The Acquisition of the Lexicon,* (433-470) Cambridge, MA:MIT Press.

Brown, Roger (1973). *A First Language.* Cambridge, MA: Harvard University Press.

Casselli, Maria; Bates, Elizabeth; Casadio, Paola; Fenson, Judi; Fenson, Larry; Sanderl, Lisa & Weir, Judy (1995). A cross-linguistic study of early lexical development. *Cognitive Development,* **10**, 159-199.

Echols, Catherine (1990). An influence of labeling on infants' attention to objects and consistency: Implications for word-referent mappings. Paper presented at the International Conference on Infant Studies, Montreal, Quebec, April, 1990.

Fisher, Cynthia; Hall, Geoffrey; Rakowitz, Susan & Gleitman, Lila (1994). When it is better to receive than give: Syntactic and conceptual constraints on vocabulary growth. *Lingua,* **92**, 333-375.

Gelman, Susan & Tardif, Twila (1998). Acquisition of nouns and verbs in Mandarin and English. In Clark, E. (Ed.), *The proceedings of the twenty-ninth annual child language research forum* (27-36). Stanford, CA, USA: Center for the Study of Language and Information.

Gentner, Dedre (1978). On relational meaning: The acquisition of verb meaning. *Child Development,* **49**, 988-998.

Gentner, Dedre (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S.A. Kuczaj, II (Ed.), *Language Development: Vol. 2 Language, Thought, and Culture* (31-53). Hillsdale, NJ: Erlbaum.

Gentner, Dedre & Boroditsky, L. (in press). Individuation, relativity, and early word learning. To appear in M. Bowerman & S. Levinson (Eds.), *Language Acquisition and Conceptual Development.* England: Cambridge University Press.

Gillette, Jane; Gleitman, Henry; Gleitman, Lila & Lederer, Ann (in press). Human simulation of vocabulary learning. To appear in *Cognition*.

Golinkoff, Roberta; Mervis, Carolyn & Hirsh-Pasek, Kathy (1994). Early object labels: The case for a developmental lexical principles framework. *Journal of Child Language,* **21**, 125-155.

Leonard, Laurence; Schwartz, Richard; Morris, Barbara & Chapman, Kathy (1981). Factors influencing early lexical acquisition: lexical orientation and phonological composition. *Child Development,* **52**, 882-887.

Macnamara, John (1972). Cognitive basis of language learning in infants. *Psychological Review,* **79**, 1-13.

Macnamara, John (1982). *Names for Things: A Study of Human Learning.* Cambridge, MA: MIT Press.

Markman, Ellen (1992) Constraints on word learning: Speculations about their nature, origins, and domain specificity. *The Minnesota symposia on child psychology, 25, Modularity and constraints in language and cognition* (59-101). Hillsdale, NJ: Erlbaum

Markman, Ellen & Wachtel, Gwyn (1988) Children's use of mutual exclusivity to constrain the meaning of words. *Cognitive Psychology,* **20**, 121-157.

Maratsos (1991). How the acquisition of nouns may be different from that of verbs. In A. Krasnegor, D. Rumbaugh, R. Schiefelbusch, & M. Studdert-Kennedy (Eds.), *Biological and Behavioral Determinants of Language Development* (67-88). Hillsdale, NJ: Erlbaum.

Nelson (1973). Structure and strategy in learning to talk. *Monographs of the Society for Research in Child Development,* **38(1-2)**.

Smiley, Patricia & Huttenlocher, Janellen (1995). Conceptual development and the child's early words for events, objects, and persons. In M. Tomasello & W. Merriman (Eds.), *Beyond Names for Things,* 21-62, Hillsdale, NJ: Erlbaum.

Smith, Linda (1995) Self-organizing processes in learning to learn words: Development is not induction. *The Minnesota Symposia on Child Psychology, 28, Basic and applied perspectives on learning, cognition, and development* (pp. 1-32). Mahwah, NJ: Erlbaum