

Semantic Constraints on Lexical Categories

William Nagy and Dedre Gentner

*Department of Psychology, University of Illinois at Urbana-Champaign,
Champaign, Illinois, U.S.A.*

Children are prodigious word learners, able to gain substantial information about a word's meaning on the basis of even a single encounter with the word in context. Two experiments are reported which test the hypothesis that the notion "possible word meaning" plays a role in constraining the inferences that adult language users make about new words encountered in context. The results indicate that subjects applied implicit knowledge of constraints on possible word meanings in making and evaluating hypotheses about unfamiliar words encountered in context and suggest that implicit knowledge of both universal and language-specific constraints on possible word meanings may contribute to the remarkable speed with which children and adults acquire new words.

INTRODUCTION

Children are remarkably efficient word learners. The average child learns 2000-3000 per year between grades 3 and 12, and some children may learn at twice that rate (Miller & Gildea, 1987; Nagy & Anderson, 1984; Nagy & Herman, 1987; Templin, 1957). Because it is implausible that more than a fraction of this word knowledge could come through explicit vocabulary instruction, it appears that children must learn substantial numbers of words incidentally, inferring meanings from written or oral context (Jenkins & Dixon, 1983).

Requests for reprints should be addressed to Dedre Gentner, Department of Psychology, University of Illinois at Urbana-Champaign, 51 Gerty Drive, Champaign, Illinois 61820, U.S.A.

We would like to thank Melissa Bowerman, Pat Smiley, Phil Johnson-Laird, and two anonymous reviewers for their insightful comments on this research.

The work upon which this paper is based was supported in part by the Office of Educational Research and Improvement under Cooperative Agreement No. OEG 0087-C1001 and by a summer research fellowship to the second author from the Max-Planck Institute in Nijmegen, Netherlands. This paper does not necessarily reflect the views of the agencies supporting the research.

How people learn word meanings from context is still poorly understood. Natural contexts are often uninformative (Beck, McKcown, & McCaslin, 1983; Schatz & Baldwin, 1986), so that gaining an adequate representation of a word's meaning should require a large number of exposures to the word in a variety of meaningful contexts (Deighton, 1959). Yet children demonstrate a surprising ability to gain quite substantial information about a word's meaning from even a single exposure to the word in a written or oral context (Carey, 1978; Dickinson, 1984; Heibeck & Markman, 1987; Nagy, Anderson, & Herman, 1987; Nagy, Herman, & Anderson, 1985). This rapid initial learning (labelled "fast mapping" by Carey, 1978) suggests that in word learning, as in reading and listening comprehension in general, the reader or listener plays an active role, bringing to bear both linguistic and world knowledge to generate hypotheses or models of meaning that go beyond the information present in the text.

This research investigates the role of linguistic knowledge and, specifically, of *semantic* knowledge, in the way adults extract meanings from context. We approach this issue from two perspectives. On the theoretical level, this work arises from basic research we have pursued independently on how language refers to the world: issues such as how different form classes refer, how referential patterns differ cross-linguistically, and how children learn word meanings (Gentner, 1975; 1978; 1981a; 1982), as well as how word reference is productively extended (Nagy, 1978). On a more practical level, this work bears on vocabulary acquisition and its relationship to reading (Nagy & Herman, 1987; Nagy et al., 1985). We hope that by understanding better the process by which readers learn new words from context, we can improve vocabulary instruction in schools.

INFERRING MEANINGS FROM CONTEXT

Consider the hypotheses one might make about the meaning of the italicised nonce word in the following text:

There, a scant two hundred yards ahead of the stopped train, he found a washed-out bridge. The whole thing had toppled into a ravine. If it had not been for the mysterious flagman, the train would have *ganted* across the ravine into the opposite embankment, killing passengers and crew.

What knowledge does a reader bring to bear in this situation? Both knowledge of the world and knowledge of one's language help limit hypotheses about the new word's meaning.

World Knowledge

First, it is clear that world knowledge plays a major role in determining what hypotheses the word-learner entertains. For example, one would not

expect *ganted* to mean something like *looked*. In this particular genre of text, trains do not look into embankments, and even if they did, that probably would not have killed the passengers and crew. Knowledge of the world permits one to construct a fairly definite scenario of the event depicted by the sentence containing the word *gant*, and this knowledge is obviously crucial in the process of making hypotheses about the meaning of the word. The problem is that world knowledge does not sufficiently constrain hypotheses about the meaning of a new word; in fact, it allows the word-learner to make inferences that increase the number of hypotheses to be considered. Table 1 lists some of the features of the scenario that a reader might construct from this text (though this list is certainly not exhaustive).

The point is that even after the application of world knowledge, there is still too much information available; there are too many hypotheses about the word's meaning that are logically consistent with the input. If the learner had to consider all the hypotheses like "cessation of linear motion in the dark resulting in loss of human life", learning words from text would hardly be possible, let alone efficient.

TABLE 1
Some Aspects of the "Train Crossing the Washed Out Bridge" Scenario

motion	not intentional
with relatively low friction	undesirable consequences
unregulated	disruption of schedule
forward	environmental damage
swift	property of value in excess of \$2000
against air resistance	impact
above water	contact between unlike substances
across ravine	increase in entropy
parallel to long axis of figure moving	non-reversible change of state
ballistic trajectory	abrupt change of shape
through the air	generation of heat
non-volitional motion of animate beings	noise
motion of contained objects relative to container	fear
great-momentum	surprise
reduction in velocity	pain
of an artefact	injury
of a vehicle	loss of life
wheeled vehicle	large objects
departure from planned route	multiple objects
disengagement of propelling mechanism	long objects
event outside of populated area	rigid objects
in the dark	sausage-like string of objects
without warning	metal

Linguistic Knowledge

General knowledge of the world can bring the learner to the point of having a fairly specific scenario, or situation model (Kintsch, 1986) associated with a piece of text containing an unknown word. The learner's task is then to discern which parts of this scenario are likely to be associated with the word's meaning. At this point, we believe linguistic knowledge comes into play: specifically, knowledge about how information is likely to be partitioned into word meanings. The work of semanticists such as Fillmore (1978), Jackendoff (1975; 1983), Langacker (1982; 1986), Miller and Johnson-Laird (1976), and Talmy (1972; 1975; 1978) suggests that there are systematic patterns of word meaning within a language. In Langacker's terms, to know what a word means, you have to know more than what situation or object it refers to; you have to know how the language construes that object or situation.

The first use of linguistic knowledge in our example is rather elementary. Our scenario has some objects—the train, the bridge, the ravine, the people—and a cluster of events. Given some knowledge of English syntax and morphology, we can deduce that *ganted* is a verb. Given the knowledge that verbs are typically associated with events and relations, and nouns with objects, we can guess that *ganted* refers to some aspects of the cluster of events, rather than to one of the participants. The association of verbs with events and nouns with objects is certainly not without exception (Maratsos & Chalkley, 1981), but probably plays an important role in restricting children's initial hypotheses about word meanings (Brown, 1957; Gleitman & Wanner, 1982; Landau & Gleitman, 1985). The syntactic context in which *ganted* occurs, and especially the preposition *into* following it, suggests that it is a verb of motion. However, this still leaves a large number of possibilities for the meaning, e.g. "motion above water", "motion of contained objects relative to container", or "motion of a group of long rigid objects".

We can restrict our hypotheses further by considering what aspects of events are typically incorporated into the meanings of English motion verbs. English motion verbs, for example, do not characteristically specify properties of the moving object. Thus, a meaning such as "motion of a group of long rigid objects", though compatible with the context, is not a likely candidate for an English verb meaning. It should be noted that rejecting this hypothesis about the meaning of *gant* is based on a knowledge of English, and not on some general notion of plausibility. The degree to which object properties can be included in the meaning of a verb stem is a parameter along which languages differ (Sapir, 1944; Talmy, 1975; 1978). As Talmy (1972) points out, in the Atsugewi language, meanings like "for a particulate substance to move or be located" are typical, rather than

exceptional. Even two closely related languages may differ in this regard; for example, Plank (1980) has argued that German differs systematically from English in that German allows verbs to place more restrictions on properties of their arguments.

We can go one step further in restricting our hypotheses about the meaning of *gant* by invoking another piece of knowledge about English semantics. Although a few English motion verbs incorporate direction (e.g. *enter, leave*), the meaning pattern "motion plus manner" is far better represented in the English lexicon (e.g. *tumble, roll, slide, swoop, dash, slink, slip, strut, saunter, gallop, hop, stroll, hurtle, skip, fly, zip, trudge*). This pattern of meaning is developed in English to a greater extent than in, for example, Romance languages (Talmy, 1972; 1975) or Japanese.

A good first guess about the meaning of *ganted*, then, is that it is a verb of motion, with some additional specification as to the manner of motion; namely, the manner in which a train would move if it went over a washed-out bridge. One can guess, then, that *gant* means something like "to move swiftly and uncontrollably", or perhaps "ballistically".

CONSTRAINTS ON WORD MEANINGS

The preceding discussion illustrates the importance of linguistic knowledge in allowing the learner to arrive at a reasonably specific hypothesis about the meaning of a new word encountered in context. The point we want to stress is that the information provided by the text and the reader's (or listener's) knowledge of the world does not constrain the range of hypotheses enough to allow rapid learning. There will always be indefinitely many hypotheses logically consistent with the data available (Quine, 1960). As Carey (1983) and Markman (1987) have argued persuasively, learning word meanings from context is possible only if there are constraints on the hypotheses that the learner makes. The nature of such constraints is still largely unexplored; but there have been a few specific proposals.

There are at least three rough kinds of biases or constraints that adult speakers might have about word meanings: (1) language-universal biases against impossible concepts; (2) language-general biases; and (3) language-specific biases reflecting knowledge of regularities, such as the English vs Spanish patterns for verbs of motion.

Keil (1979; 1981) and Sommers (1963) have postulated a particular type of conceptual constraint on word meanings: ontological constraints on natural kind terms. Keil (1983) has found that even young children have some ability to use knowledge of ontological categories in making hypotheses about the meanings of new words; for example, if a child is told that "Throstles can be fixed", the child can generally correctly answer a question such as "Can throstles bloom?" Keil's stronger claim, that ontolo-

gical categories must form a strictly taxonomic hierarchy (the so-called M-constraint on ontological categories), has been disputed by Carey (1983). She suggests that natural kind concepts are constrained by children's emerging theories of the world, rather than by purely structural knowledge. But both accounts agree on some kind of constraints on concepts that would guide inferences about new words.

The constraints or regularities discussed by Keil and Sommers could be thought of as reflecting possible concepts that a word might refer to. We turn now to proposals concerning linguistic constraints, i.e. constraints on which kinds of possible concepts are acceptable as word meanings. The research bearing most directly on the psychological reality of constraints or biases on possible word meanings has been in the area of early language acquisition. One of the first and most basic semantic constraints to be investigated was the *one-to-one mapping principle* (Slobin, 1973). There are several other variations of this hypothesis—the contrast principle (Clark, 1983), the mutual exclusivity hypothesis (Markman, 1987), and the one-to-one mapping hypothesis (Pinker, 1984). Although these proposals differ in detail, all postulate some sort of expectation on the part of the language learner that there will be a one-to-one mapping between linguistic elements and conceptual elements.

Another basic constraint on meanings that appears to apply early in language acquisition is the *taxonomic constraint*. Markman and Hutchinson (1984) asked children aged 2–5 years old to choose which of two pictures went with a third target picture under either neutral or linguistic instructions. The choice was always between a thematic and taxonomic associate. For example, children were given a target picture of a cow, and asked to decide whether it went with *milk* (a thematic associate) or with a *pig* (a taxonomic associate). In this neutral or non-linguistic condition, children were equally likely to choose the thematic or the taxonomic associate. However, in the linguistic version of the sorting task, children were told, for example, "See this fep? Put it with another fep." In this case, children sorted predominantly on a taxonomic basis, putting the cow with the pig, rather than with the milk (see also Hutchinson, 1985). In the non-linguistic condition, then, children find thematic relations at least as salient as taxonomic relations; but when the task requires thinking in terms of the meaning of a word, children focus on taxonomic relationships. This is presumably helpful in word learning, because noun categories in English (and quite possibly in all languages) tend to be organised taxonomically. There may be a strong association between *cow* and *milk*, but English does not have words with a meaning like "cow and milk" or "cow or milk". Thus, even very young children understand some of the principles that constrain word meanings over and above any constraints on conceptual association.

The one-to-one mapping hypothesis and the taxonomic constraint are not the only examples of possible language-general biases or constraints. Another example (many instances of which are tested in Experiment 1) is biases against lexicalising highly contingent concepts. There are many possible concepts that one can easily entertain but for which one would not wish to dedicate a particular word meaning, by reason of their specificity or ephemerality. For example, it would be surprising to come across a verb that meant "to drive a 1984 Buick at 10 a.m. in Kansas City", although this is a perfectly possible concept. We postulate a general bias against the proliferation of highly contingent word meanings; people's semantic knowledge¹ includes a knowledge of which kinds of meanings are liable to be conventionally useful. It must be noted that many languages permit highly specific meanings to be captured with morphologically transparent derivations. Thus, we could use a denominal verb like *Buicked* to denote driving a Buick, or *teapotted* to denote hitting someone with a teapot (Clark & Clark, 1979). As long as the neologisms are morphologically transparent, their meanings can be derived from existing meanings; hence they do not constitute cases of lexical proliferation.

The clearest examples of semantic, as opposed to purely conceptual, constraints are those that are language-specific. Several linguists have suggested that languages can differ in the preferred semantic patterns they use within a given domain (Casad & Langacker, 1982; Fillmore, 1978; Lakoff, 1987; Langacker, 1982; Talmy, 1972; 1975). For example, Talmy (1972; 1975) suggested that there are language-specific semantic patterns for verbs of motion. All verbs of motion (by definition) have change-of-location as part of their meaning. However, languages differ in which other semantic elements they include in their motion verbs (or, in Talmy's terminology, in which other elements are conflated into the verb's meaning). In English, we readily conflate *manner of motion* into our motion verbs (as in *float*, *hop*, *stride*, and so on). In contrast, Latin languages such as Spanish and French almost never conflate manner of motion into the verb, but often conflate the *path of motion with respect to some ground* (e.g. French *entrer*, *salir*, *descendre*, and so on). Talmy gives the example of a bottle bobbing up and down in a river and moving into a cave. In English we would say "The bottle is floating into the cave", conflating the floating manner in with the change of location and leaving the path ("into the cave") in a separate prepositional phrase. In Spanish, the likely description would be "La botella entro a la cueva flotando." Here the path

¹Some might prefer to call this pragmatic knowledge. While agreeing that this kind of knowledge has its pragmatic aspects, we prefer the term *semantic* because we are talking of knowledge of the likely meaning of words.

of motion is conflated into the main verb and the manner is placed in a separate adverbial/participle.

There is little experimental evidence on whether language-specific patterns of lexical semantics constrain the hypotheses that learners make about the meanings of new words. However, there is evidence that children recognise language-specific semantic regularities among word meanings in ways that influence their acquisition and use of new words. A variety of evidence along these lines comes from the work of Bowerman (1981; 1982a; 1982b) on children's late semantic errors. These errors occur relatively late in language acquisition, after the child has used the terms correctly for a long period, and seem to arise, paradoxically, from the child's increasing understanding of the dominant underlying semantic patterns of the language. A typical instance concerns the expression of causality. At about the age of 3–7 years, children begin to say things like "Who highered the swing?" (boy, 5:11); "Don't dead him" (as M. picks up a spider) (Eva, 4:10); or "I'll jump that down" (Rachel, 4:9)—meaning "cause that to go down by jumping on it" (R. is about to jump on a bathmat placed on top of the water in a tub) (Bowerman, 1982b, pp. 14–18). These late errors seem to show the child's increasing sense of the semantic structure of the language. The child seems to have grasped that, in English, the verb for causing a particular change of state is often the same word as the verb for the change of state, or even the word for the state itself. For example, the same word (*open*) is used in "Open the door" (i.e. cause the door to become open), "The door opens" (the door is changing to a state of being open), and "The door is open" (the door is in the state of being open). According to Bowerman, such regularities may go unappreciated at the earliest stages of semantic acquisition; but as the children acquire expertise in English, they implicitly grasp semantic regularities, as evidenced by their attempts to generate utterances based on this pattern.

For our purposes, Bowerman's research is important in showing that children have knowledge of patterns of permissible word meanings and also of morphological regularities among sets of word meanings. Clark (1982) and Clark and Hecht (1982) document other instances of children creating new words according to regularities they have discovered. For example, Clark (1981) notes that 3-year-olds spontaneously use the *-er* suffix to indicate agency, as in *crayoner* (one who uses crayons) or *rainer* (one who drives rain away). Such errors are clearly self-generated, not learned by imitation; and although the particular utterances are erroneous, they show that the child is beginning to grasp a semantic pattern that plays a role in an adult's competence in English. Such errors are therefore further evidence of naturally developing knowledge of what constitutes allowable word meanings in one's language.

Gentner (1981a; 1982) has discussed another way in which the learning of language-specific semantic patterns may affect children's language acquisition. She suggested that verbs and other relational terms vary more in their semantic patterns, cross-linguistically, than do object-reference terms, and further that this difference affects children's acquisition of word meaning. To acquire noun meanings, the child need only realise that nouns are names for cohesive separable objects; but to learn verb meanings, the child must also learn how relations are lexically partitioned in his or her language. Consistent with this reasoning, it has been found that children's early vocabularies across many different languages contain many more nouns than relational terms (Gentner, 1982; Goldin-Meadow, Seligman, & Gelman, 1976; Nelson, 1973).

THE PSYCHOLOGICAL STATUS OF CONSTRAINTS

There is little doubt that word learning could hardly be efficient, or even possible, if learners had to consider all the hypothetical word meanings that were logically consistent with the context in which the word was encountered. However, the psychological status of the constraints that have been postulated by Keil, Sommers, Markman, and others has been subject to debate. Nelson (1988) has argued that the notion of linguistic constraints on possible word meanings is inadequate, and that the efficiency of word learning is better explained on other grounds. In particular, she argues that Markman and others have not provided evidence that their hypothesised constraints are universal, innate, specifically linguistic, or absolute constraints (rather than just biases or preferences).

We (and, we suspect, some others who have used this term) would readily agree with Nelson on at least some points—that there is still far from sufficient evidence to support any claim of innateness, and that the current evidence appears to reflect biases, rather than absolute constraints. We would not agree, on the other hand, with Nelson's apparent claim that abandoning innateness and absoluteness strips the notion of constraints of any explanatory force. Nor do we believe that the recognition that lexical development is in part a social convergence process obviates the need for linguistic and, specifically, semantic knowledge in word learning.

In the present research, we set aside questions of innateness and focus on the extent to which there are linguistic constraints or biases—both language-general and language-specific—on the hypotheses adult learners make about the meanings of new words. We will continue to use the word *constraints*, although we acknowledge the problems with this word that Nelson has pointed out, simply because we can think of no more felicitous term for “knowledge that tends to restrict the range of hypotheses consi-

dered". We present two studies that bear on different aspects of this problem: Both are exploratory, yet both produce encouraging results and may serve to suggest useful methodologies for studying semantic knowledge. In the first study, we tested a wide variety of possible semantic biases to discover which ones people are sensitive to. In the second study, we focused on two constraints for which there is fairly good linguistic evidence and asked whether people are influenced by these regularities in forming implicit word meanings from context.

If we are ever to go beyond initial demonstrations of the existence of constraints, we need a better understanding of the scope of the phenomenon. From what has been said so far, it is clear that the number of semantic regularities that a speaker might apprehend is potentially rather large. Yet in the literature to date, only a small number of constraints has been explored. Therefore, the first study was designed to explore a wide variety of possible constraints. The idea was to compare the hypotheses people make concerning the meaning of an unfamiliar word with those they make when given a similar non-linguistic situation. The subjects were shown a context sentence containing either a blank or an unfamiliar word, e.g. "Jack _____ into the canyon" or "Jack strounded into the canyon". In both cases, the subjects were given a set of possible concepts and asked to rate them either as to how well they fit into the blank (in the non-linguistic or cloze condition), or as to how plausible they were as meanings for the unfamiliar word (in the word-meaning condition). This task, with its contrast between linguistic and non-linguistic contexts, is inspired by the research of Markman and Gelman and their colleagues, as well as by the pioneering studies of Werner and Kaplan (1952). The set of possible semantic regularities investigated in this experiment is, frankly, a rather motley one. We included a few constraints for which there exists good psychological evidence, such as Markman's taxonomic constraint on nouns and Talmy's manner-conflation in English motion verbs. We also included a large number (15 in all) of possible biases designed by ourselves to be roughly parallel to those suggested by linguists for other languages. Many of these were variations of the "no lexical proliferation" principle.

The four alternatives given to the subject varied in their semantic acceptability and also in their conceptual plausibility in a 2×2 design (as shown in Table 3). If adult subjects are sensitive to semantic, as opposed to purely conceptual, regularities, then we expect the following pattern: In the non-linguistic condition, subjects will prefer the two conceptually plausible alternatives over the two conceptually implausible alternatives, but will show no preference for semantic plausibility. In the word-meaning condition, subjects will continue to be sensitive to conceptual plausibility but will also make a further distinction; they will prefer the semantically plausible over the semantically implausible alternative. Thus, if people

possess specifically semantic knowledge in addition to their conceptual knowledge about word meanings, then we should see an interaction between condition (linguistic or non-linguistic) and type of regularity (semantic or conceptual). If, on the other hand, subjects fail to make a distinction between semantic and conceptual plausibility, then we should see the same pattern of preference across linguistic and non-linguistic tasks.

EXPERIMENT 1

The Constraints

Several classes of constraints were investigated, some purely semantic and others partly semantic and partly morphological. Some of these constraints seem likely to reflect universal tendencies, whereas others appear to be language-specific restrictions. As mentioned above, the study of constraints is still in its infancy, and therefore this assignment is tentative.

Table 2 gives examples of possible constraints examined in this experiment. The first constraint listed is the *taxonomic* constraint on nouns,

TABLE 2
Examples of Constraints Tested in Experiment 1

Nouns

1. Taxonomic constraint: Nouns tend to reflect taxonomic rather than thematic relationships:
needles and pins
*needles and thread
2. Durative constraint: Nouns are usually differentiated in terms of long-term properties rather than short-term properties:
red bird
*flying bird (i.e. bird on the wing)

Verbs

3. Time of day constraint: English verbs do not include reference to specific times of day:
to wake up quickly
*to wake up early
 4. Cessation constraint: English verbs do not generally have meanings of the form "to stop Xing" for any *specific* meaning X:
to laugh nervously
*to stop sewing
-

Note: Phrases marked with an asterisk are meanings that are *not* likely to be expressed by a single word in English.

which states that nouns tend to refer to classes, not to thematic collections. Markman and Hutchinson's (1984) research indicates that young children observe this constraint, and indeed it seems likely to be a language-universal constraint. The second constraint, the *durative* constraint on nouns, which states that nouns tend to refer to classes characterised by long-term properties, again seems likely to reflect a universal tendency, but it admits some exceptions, especially in the case of nouns morphologically derived from verbs (e.g. *explosion*, *recipient*). We next turn to two constraints that seem less likely to be universal. The next constraint is the *no specification of time of day constraint* on verbs. It states that although verbs in English are inflected for tense, they do not incorporate in their meanings any specification of the time (in terms of a clock or calendar) at which an activity is performed. Hence, we would predict that one is much less likely to postulate a new verb meaning "to wake up early in the morning" than one meaning "to wake up quickly".² The last constraint in Table 2, the *no cessation of X constraint* on verbs, states that there are no English verbs meaning "to stop doing a particular action". Thus, although there are general negative verbs like *stop*, *cease*, *halt*, and *desist*, they mean simply to stop doing whatever one is doing. We do not find verbs meaning "to stop writing" or "to stop running".³ A discussion of all 15 of the constraints tested in this experiment is given in Appendix 1.

Method

Subjects. The subjects were 68 undergraduate students at a large midwestern university, who volunteered to participate in the experiment. They were divided randomly into two groups, the *cloze* (non-linguistic) group and the *definition* (linguistic) group.

²Note that verbs like *to winter*, *to summer*, or *to dawn*, which do incorporate time, do so without specifying any other activity. Thus, *to winter* can mean "to spend the winter", but not something like "to engage in winter sport". More serious counterexamples are *to lurch* or *to dine*. However, all these examples are morphologically transparent derivations from nouns, rather than independent verbs.

³Apparent counterexamples to the "no cessation of X" constraint are *to halt*, *to still*, *to calm*, and *to give up*. These do not constitute exceptions because they are general cessation verbs; they do not specify the action which is stopped. Two other possible counterexamples are *to freeze* (to stop molecular motion) and *to wake up* (to stop sleeping). However, these events seem likely to be intuitively perceived as positive (i.e. "to become cold and form ice" and "to open one's eyes, increase one's activity, etc.", respectively). A more serious counterexample is *to resign* (to stop serving in some capacity), which does appear to be a genuine exception to the "no cessation of X" rule. The existence of such exceptions makes it clear that the regularities we are discussing are biases, not absolute strictures.

Materials. The experimental booklets contained 20 items, all in either the cloze or the definition condition, as shown in Table 3, along with instructions appropriate for that condition. In the cloze condition, the sentence contained a blank, followed by four phrases. In the definition condition, the item consisted of the same sentence, with a nonce word in place of the blank, and the same four phrases. Note that the content of the two conditions is identical; the difference is in whether it contains a blank or a nonce word.

For each item, the four phrases differed orthogonally on two dimensions: how well they fit the context ("contextually appropriate/inappropriate"), and whether they are allowable word meanings in English ("lexically appropriate/inappropriate"). For example, in the item given in Table 3, both "laugh nervously" and "stop sewing" fit the context, but only "laugh nervously" represents a possible word meaning in English because "stop sewing" violates the *no cessation of X* constraint. The contextually inappropriate phrases are "spend time on the beach" and "arrive exactly on the hour". Of these two phrases, "arrive exactly on the hour" is less likely to constitute a possible word meaning in English than "spend time on the beach", because it violates the *no specification of time* constraint.

Procedure. Each subject received a booklet of 20 items, all either in the cloze or in the definition condition. Each item consisted of a sentence

TABLE 3
Sample Item from Experiment 1 in the Two Experimental Conditions

Cloze condition

Martha would _____ whenever she heard footsteps on the sidewalk in front of her house.

- | | |
|-----------------------------|--|
| (a) spend time on the beach | |
| (b) laugh nervously | |
| (c) refuse to agree | |
| (d) stop sewing | |

Definition condition

Martha would werpet whenever she heard footsteps on the sidewalk in front of her house.

- | | |
|-------------------------------------|--|
| werpet: (a) spend time on the beach | |
| (b) laugh nervously | |
| (c) refuse to agree | |
| (d) stop sewing | |
-

- Key:* (a) spend time on the beach—contextually inappropriate, lexically appropriate.
 (b) laugh nervously—contextually appropriate, lexically appropriate.
 (c) refuse to agree—contextually inappropriate, lexically inappropriate.
 (d) stop sewing—contextually appropriate, lexically inappropriate.

followed by four phrases, as described above. In both conditions, the subjects were to rate each phrase independently. In the cloze condition, the subjects were instructed to rate each phrase as to how well it fit into the blank, on a scale of 1–5, from “utterly implausible, unnatural” to “perfectly natural, totally plausible”. In the definition condition, the same phrases were rated on a similar plausibility scale, but in this case the subjects rated how plausible the phrases were as possible meanings for the underlined nonce word. A single random order of items was used in all booklets. The subjects were self-paced.

Design and Analysis. A $2 \times 2 \times 2$ mixed factorial design was used, with task condition (cloze or definition) as a between-subject factor, and lexical appropriateness and contextual appropriateness as within-subject factors. We expected an effect of contextual appropriateness in both tasks; but, more importantly, we predicted that lexically inappropriate phrases would be rated lower in the definition condition than in the cloze condition. Thus, we expected a significant interaction between task and lexical appropriateness. Such an interaction would indicate that there are constraints on what sorts of information people can package into word-meanings—constraints that are distinct from the concept of “plausibility in context”.

Results and Discussion

Figure 1 shows subjects' mean plausibility ratings. It can be seen that, as predicted, both groups of subjects preferred contextually plausible meanings; but, more importantly, the definition group showed a further bias against lexically implausible meanings. A mixed-measure analysis of variance across subjects was performed for the $2 \times 2 \times 2$ design of task \times contextual appropriateness \times lexical appropriateness.

We begin by describing the main effects. (Although the main effects are not of primary interest, they serve as a test of whether the materials worked as planned.) The main effect of task (cloze vs definition) was not significant [$F(1,66) = 1.47, p > 0.2$], confirming that items were not rated higher in one condition than in the other. The effect of contextual appropriateness was highly significant [$F(1,66) = 970.9, p < 0.001$]; items which were intended to fit the context well were rated consistently higher than items which did not fit the context. The main effect of lexical appropriateness was not significant [$F(1,66) < 1, p > 0.5$]. That is, overall, the lexically inappropriate phrases were not rated as being more or less plausible than the lexically appropriate phrases.

The predicted interaction of lexical appropriateness with task was highly significant [$F(1,66) = 17.4, p < 0.001$]. Lexically inappropriate phrases—phrases that should not constitute possible word meanings in English—were rated considerably lower in the definition condition than in the cloze condition.




FIG. 1. Results of Experiment 1: Mean ratings of naturalness in context for items in the cloze and definition conditions.

There was also a three-way interaction, reflecting the fact that the interaction of condition and lexical appropriateness was found only for those items that fit the context. As expected, items that were contextually inappropriate showed a floor effect; all such items received low ratings, regardless of task or lexical appropriateness.

In addition to the subjects analysis, an items analysis was performed. The main effects were the same, with the exception that the effect of task was significant in the items analysis [$F(1,19) = 7.5, p = 0.013$]. The key interaction of lexical appropriateness \times task remained significant [$F(1,19) = 12.6, p = 0.002$]. Hence, this interaction is to be generalisable both across subjects and across items [quasi min $F(1,49) = 7.3, p < 0.01$].

One aspect of the data appears problematic. As Fig. 1 shows, the ratings for lexically inappropriate phrases are *higher* than those for lexically appropriate phrases in the cloze task and *equal* to those for lexically appropriate phrases in the definition task. Our hypothesis would suggest

that these ratings should be *equal* in the cloze task, and that lexically appropriate phrases should be rated *higher* than lexically inappropriate phrases in the definition condition. We believe that the explanation for this apparent anomaly lies in our method of stimulus construction. In order to avoid biasing the materials in favour of our hypothesis, we took great pains to make the contextually appropriate but lexically inappropriate phrases fit the context at least as well as the contextually lexically appropriate phrases. We appear to have succeeded all too well. In the cloze task, in which the two contextually appropriate choices should have been equally good, the lexically inappropriate phrase is rated as far superior to the lexically appropriate phrase. In the definition task, this difference disappears, as though the contextual superiority of the CA-LI items was balanced by the lexical superiority of the CA-LA items. The crucial result in this experiment, however, does not have to do with the absolute level of the ratings, but with the *difference* between the two tasks. Only one category of phrases—the contextually appropriate but lexically inappropriate items—shows a difference in ratings between the two conditions. This indicates that subjects evidenced sensitivity to *semantic* constraints *only* when evaluating different hypotheses about the possible meaning of an unfamiliar word. These results provide clear evidence of a distinction between conceptual-contextual plausibility and lexical plausibility.

The items in this experiment were intended to include both universal and language-specific constraints on possible word meanings. All of the six items that represented what are likely to be universal constraints on the form of word meanings, or at least universal tendencies, showed the predicted pattern of results. Of the 14 items that represent what we judged to be language-specific constraints—restrictions on what is a possible word in English, which may not apply in other languages—all but three items behaved consistently with our predictions. Thus, it appears that not only do people approach language learning with language-general expectations about possible word meanings (e.g. that nouns are likely to represent taxonomically rather than thematically organised categories), but that language learners acquire implicit knowledge of the semantic patterns of their language, and use this knowledge when formulating hypotheses about the meanings of new words.

The results of Experiment 1 are promising. They indicate that speakers possess and use a variety of constraints on possible word meanings, including some that are apparently language-specific. However, one concern with the method used in Experiment 1 is that it relies on subjects' explicit judgements of possible word meanings. Thus it was possible that subjects were using formal knowledge about word meanings that they would not use in ordinary processing. The second experiment avoids this difficulty. Instead of directly asking subjects to judge possible word mean-

ings, we used a naturalistic task resembling learning from context. The subjects encountered a new word, first in a rich context and then in a neutral context, and were later asked what inference they derived from the second use. The key question was whether the kind of linguistic cues present during the initial use of the word would influence the inferences subjects made on encountering a subsequent use of the word.

Experiment 2 differs from Experiment 1 in another way. In Experiment 1 we considered a wide variety of constraints on word meaning; in Experiment 2 we focus on two well-documented constraints governing the possible meanings of nouns and verbs in English. As Talmy (1972; 1975) has discussed, English verbs frequently incorporate information about the instrument used to perform an action or the manner in which an action is performed, but rarely incorporate information about the properties of the object upon which the action is performed. For example, English does not tend to have many meanings like "to carry something long and rigid", or "to find something small". This restriction on verb meanings is not universal among languages. There are languages, for example, which have different verb stems for the meanings "to carry a long rigid object" and "to carry a small round object" (Friedrich, 1970; Talmy, 1972; 1975). In fact, Plank (1980) argues that German systematically differs from English in its somewhat greater tendency to allow verbs to specify object properties, e.g. *schliessen* (to shoot a game animal or bird) vs *erschliessen* (to shoot a human or a non-game animal, particularly one to which one is attached). Therefore, the first constraint we considered in Experiment 2 was that English verbs incorporate manner but do not incorporate object properties.⁴

⁴ An apparent counterexample to the "no incorporation of patients in verbs" principle is given by Clark and Clark (1979), who document examples of verbs such as *paint* and *cover*, which seem to involve object incorporation ("to put X on" something). However, as Clark and Clark point out, these verbs result from a specific morphological process of denominalisation and are marked by their transparent phonological relationship with the nouns from which they are derived. There are some verbs such as *to lug*, which might be paraphrased "to carry something heavy". However, it could be argued that it is still the manner in which the action is performed, and not the property of the object, that is essential to the meaning of *lug*. No matter how heavy the object, if someone is strong enough to carry it with ease, he or she would not be said to be lugging it. Other possible counterexamples are verbs like *to shatter*, *to spray*, *to tear*, and *to fold*, all of which require objects with certain specifications. However, it could be argued that these verbs simply specify a particular manner of change of state, and that the object specification simply follows from the manner, e.g. any object that can be folded will do as the object of *fold*. Weather verbs, such as *to rain* and *to snow*, are well known counterexamples to the no-patient-incorporation rule, but they occur only in a restricted semantic domain. More serious counterexamples are *to kill* vs *to murder* vs *to assassinate*. The latter two require a human patient, and the last requires a politically prominent human. (We thank the anonymous reviewers for pointing out this example as well as other problematic cases.) This is a clear instance of patient-incorporation, which cannot be

The second constraint concerns nouns. English nouns that are object names do not typically incorporate information about actions that have been performed on their referents. Again, there are morphologically transparent instances such as *recipient* or *nominee*, with meanings of the form "person who has been Xed" for some verb *X*. There are also more serious exceptions, such as *victim* (although it could be argued that *victim* is a rather general term for "one who has incurred harm" and does not specify the particular action that has been performed). One does not expect to find, at least in English, monomorphemic terms with meanings like "fish caught by nets (as opposed to hooks)" or "a long stick being waved in large circles".

Prior research on whether there are lexical constraints specific to particular parts of speech has produced mixed results. Brown (1957) showed preschool children pictures of novel actions being performed on novel substances in novel containers, used a novel word in a sentence to describe the picture (e.g. "In this picture you can see a niss"), and then asked the children to choose which of several other pictures also depicted the concept referred to. He found that children's choices depended on the grammatical context of the word. If it was used as a verb, children associated it with the action; if it was used as a count noun, they associated it with the container; and if it was used as a mass noun, they associated it with the substance. In a similar study, Katz, Baker, and MacNamara (1974) studied the extent to which 2-year-olds use the common noun-proper noun distinction (also signalled by the presence or absence of articles) to constrain hypotheses about the meanings of new words. When they showed 2-year-olds a new doll and said "This is Dax", the 2-year-old girls (but not the boys) would usually take *Dax* to be a proper name for that doll, and not apply it to other dolls. If, on the other hand, they said "This is a dax", the children were more likely to take *dax* to be a common noun, and apply it to other similar dolls. Similarly, Wykes and Johnson-Laird (1977) found that 3-year-olds learn something about the selectional restrictions of verbs from hearing them used in context several times. For example, after several exposures to

explained as manner-specification. However, according to Plank (1980), English is considerably less prone to allow such patient incorporation than is German. We have mentioned German's two verbs for shooting, *schliessen* and *erschliessen*. There are also *abstechen* (animals) vs *erstechen* and *niederstechen* (people) as compared to the English *to stab*; and *ersaufen* (animals) vs *ertranken* (people) as compared to the English *to drown*. German, like English, prefers to restrict *to massacre* (*massakrieren*) to humans; but it is again more selective than English in separating *to slaughter* into an animal version (*schlachten*) and a human version (*abschlachten*), etc. The contrast is even more apparent when we compare English with American Indian languages such as Atsugewi, which makes frequent use of patient incorporation. Atsugewi has verbs which convey meanings roughly like "it dirtied into the river" or "it small-round-thinged into the river" (Talmy, 1972).

a new verb whose subject was always a liquid, the children correctly chose another liquid as the proper subject of the verb 46% of the time.

However, Soja, Carey, and Spelke (1985) failed to find evidence that 2-year-olds can use the syntactic distinction between mass and count nouns in deriving word meanings. They used a triads task in which children were shown either an object or a pile of some substance, with or without syntactic clues such as "this is a blick" or "This is some blick". Then the children had to classify the object or substance with either another object of the same form but different substance or else with objects of the same substance but different form. The results showed that although the children tended to honour the distinction between object and substance—classifying objects with objects of the same form and substances with like substances—their tendency to do so was not increased by the addition of the linguistic correlate of mass vs count noun.

Thus, the evidence so far is somewhat inconclusive with respect to form-class specific constraints on word meaning. However, for our purposes, the use of constraints specific to form class has certain advantages. First, it allows us to circumvent the problem of implausible concepts. That is, one way to be sure that our results stem from linguistic constraints, as opposed to simple conceptual-plausibility factors, is to use concepts that are perfectly plausible in all cases, but that are inappropriate to either a noun category or a verb category. A further advantage of the particular constraints tested here is that they seem likely to be language-specific, as discussed above. If these inferences show that subjects' initial meaning was influenced by the form class of the word, then this will provide evidence for the use of semantic constraints in learning word meanings from context.

EXPERIMENT 2

In Experiment 2 we tested whether English-speaking adults would make use of semantic regularities when learning the meanings of words from context. We utilised a task intended to resemble natural word learning during reading. A nonce word appears twice in a story. The first instance occurs in a rich context that allows inferences both about the manner in which an action is performed and about the object upon which it was performed. The second instance of the nonce word occurs in a context vague with respect to both types of information. The subjects are then asked to make inferences about the event described in the second context. The point of the study is that the second context provides few if any clues as to the word's meaning. Therefore, the subject inferences can be used to deduce the meaning that the subject had implicitly derived for the nonce word on the basis of the first exposure. This method has the advantage that subjects do not have to introspect about possible meanings; they merely

have to use a newly learned word in comprehending the second passage—a situation that often occurs in natural word learning.

The key manipulation is whether the nonce word is a noun or verb. Our prediction is that as we vary part of speech, different types of information supplied by the first context should be incorporated into the meaning of the nonce word. This will, in turn, lead subjects to make different inferences when they read the second occurrence of that word. This experiment might be likened to dipping a magnet into a mixture of iron filings and sand: The iron filings should stick and the sand should fall off. In this study, the picture is slightly more complex, in that we have two kinds of magnets. When the target word is a noun, we expect certain pieces of the scenario to stick to it; namely, properties of the objects referred to. When the target word is a verb, we expect a different set of pieces to stick to it; namely, information about the manner in which the action was performed. By examining the subjects' inferences in the second, underspecified context, we can see whether this differential selection hypothesis is correct.

Method

Subjects. The sample comprised 56 undergraduates at a large midwestern university, who were paid for their participation in the experiment.

Materials. There were four experimental passages, as well as one practice passage, each containing two instances of a target nonce word. The first instance of the target word was in a rich context. That is, the preceding and surrounding context were designed to give the reader a clear picture of the events described by the sentence—both the nature of the object which was being acted upon, and the manner in which the action was performed. Crucial sections from one experimental passage are shown in Table 4 (the entire passage is given in Appendix 2). In the noun version of the text, for example, it is clear from Context 1 that the noun *saptyn* refers to a small animal, probably the edible, rodent-like animal already referred to. The context also makes the nature of the actions clear; in this case, capturing the animal involves shooting seeds through a peashooter.

The second instance of the target word is in a context which is intended to be vague about both the manner of the action and the object. In the text in Table 4 (noun version), for example, when *saptyn* occurs in Context 2, there is no way to tell from the context either what kind of animal was involved, or the exact manner of capture that was intended; any such information must be supplied by the target word itself.

There were two versions of each experimental passage—a noun version and a verb version—which differed in whether the nonce word (e.g. *saptyn*) appeared as a noun or a verb (see Table 4). In the noun version,

TABLE 4
Key Sections from the Story "The Darsts of Zurdine"

Noun version

... Their way of life depends heavily on the use of the seeds of the Yrang plant, which begin to sprout between one and two seconds after they have been exposed to any moisture, and grow at speeds of up to 15 feet a second. ... Our goal that day was to bring in a particular species of smaller rodent-like animals which are prized as a special delicacy by the Darsts. ... Finally, after several hours, we caught sight of a small animal moving to our right. While the rest of the party stood motionless, the head of the hunting party pulled out a hollow reedlike peashooter, and moistened some Yrang seeds with saliva. She slipped them into the peashooter and with accuracy that amazed me, she ensnared the unsuspecting *saptyn* ...

... One of the younger Darsts, who move even more slowly than their languid elders, was somewhat behind the group, and not even making much effort to stay close. Without warning, the young Darst found himself face to face with a *saptyn*. The youngster tried to ensnare it, but in his inexperience he dropped several moistened seeds. By the time I got to him, a vine was wrapped chokingly tight around his neck ...

Verb version

... Their way of life depends heavily on the use of the seeds of the Yrang plant, which begin to sprout between one and two seconds after they have been exposed to any moisture, and grow at speeds of up to 15 feet a second. ... Our goal that day was to bring in a particular species of smaller rodent-like animals which are prized as a special delicacy by the Darsts. ... Finally, after several hours, we caught sight of a small animal moving to our right. While the rest of the party stood motionless, the head of the hunting party pulled out a hollow reedlike peashooter, and moistened some Yrang seeds with saliva. She slipped them into the peashooter and with accuracy that amazed me, she *saptyned* the unsuspecting animal ...

... One of the younger Darsts, who move even more slowly than their languid elders, was somewhat behind the group, and not even making much effort to stay close. Without warning, the young Darst found himself face to face with an animal. The youngster tried to *saptyn* it, but in his inexperience he dropped several moistened seeds. By the time I got to him, a vine was wrapped chokingly tight around his neck.

Note: Target words are italicised here for clarity; they were not italicised in the text given to the subjects.

the position of the verb in the verb version is taken by a more generic verb (in this example, *ensnare*), which does not provide crucial information about the manner in which the action was performed. Conversely, in the verb version, the position of the noun in the noun version is taken by a more general noun (in this case, *animal*), which does not provide crucial information about the nature of the object. For each passage, four questions were constructed: two about the first context and two about the second context. For each context, one question was about the manner in which the action was performed, and the other about the properties of the object. Table 5 shows the questions used for the passages in Table 4.

TABLE 5
 Questions for the Sample Passage from Experiment 2

Questions concerning the first context containing the target word:

1a. How did the head of the hunting party ensnare the animal?

(Circle a letter)

- (a) by throwing a net over it;
- (b) by shooting moistened Yrang seeds around it with a peashooter;
- (c) by throwing moistened Yrang seeds around it;
- (d) cannot be determined from the text.

1b. What sort of animal did the head of the hunting party catch?

(Circle a letter)

- (a) a large predator;
- (b) a small rodent-like, edible animal;
- (c) a poisonous razor-spined herbivore;
- (d) cannot be determined from the text.

Questions concerning the second context containing the target word:

2a. What sort of animal did the young Darst try to catch?

(Circle a letter)

- (a) a large predator;
- (b) a small rodent-like, edible animal;
- (c) a poisonous razor-spined herbivore;
- (d) cannot be determined from the text.

2b. How did the young Darst try to ensnare the animal?

(Circle a letter)

- (a) by throwing a net over it;
 - (b) by shooting moistened Yrang seeds around it with a peashooter;
 - (c) by throwing moistened Yrang seeds around it;
 - (d) cannot be determined from the text.
-

Note: Each subject saw only one question for each context. Question types were counterbalanced, so that if the subject saw the question about object properties for the first context, the second question dealt with manner of action, and vice versa. For all four of these questions, the carry-over answer (i.e. the answer reflecting the information in the first context in which the target word occurs) is choice (b).

Our primary interest is in subjects' answers to the questions concerning the second context. At issue is whether subjects are likely to carry over information about the *manner* of the young hunter's action (i.e. whether or not he had attempted to use a peashooter), information about the kind of *objects* acted upon (i.e. whether the animal was a small, edible rodent, or some other kind of animal), or both. The context of the second occurrence of the target word provides no conclusive information about either. Moreover, the noun and verb versions differ only in the part of speech of the target word, and not in the information provided by the context.

Therefore, if subjects assign different meanings to the target word, it must be on the basis of the first context. Our prediction was that subjects would retain different contextual features from the first occurrence of the word, depending on whether it appeared as a noun or a verb. Thus, subjects would respond differently to questions concerning Context 2; those reading the noun version would associate object properties with the target word, whereas those reading the verb version would associate information about the manner in which the action was performed. In addition to the key questions about Context 2, the subjects received similar questions about the first context to ensure that they had understood the materials. For questions regarding the first context, we expected subjects to give answers reflecting both types of information (manner of action and kind of object) provided by the context.

Each subject saw two questions for each story—one about the first context, and one about the second. One of these concerned the properties of the object and the other the manner in which the action was performed. There were four experimental passages and one warm-up passage. Conditions were counterbalanced so that all questions occurred equally often across subjects.

Procedure. The subjects were given a booklet containing instructions, one practice passage, and the four experimental passages. They were told that the experiment concerned how people make inferences during normal reading, and that they should therefore approach this task in the way they would approach an everyday reading situation. They were also told that the texts contained one or two words they might not have seen before, that they could proceed at their own pace, and that they were allowed to look back at the text while answering questions about it.

Design. A $2 \times 2 \times 2$ design was used, with context order (first or second), part of speech (noun or verb passage version), and question type (manner of action or object properties) as within-subject factors.

Results

Figure 2 shows the mean proportion of subjects giving carry-over answers to the questions, i.e. answers reflecting the information provided by the first context. If subjects were governed solely by *conceptual* carry-over from the first context, they would select carry-over answers independently of the form class of the target word. If, however, subjects are sensitive to the form class of the target word, they will show selective carry-over. When the target word is a noun, they will give carry-over answers to object-property questions, but not to manner-of-action questions. When the target word is a verb, they will give carry-over answers for manner-of-




FIG 2. Results of Experiment 2: Percentage of carry-over answers.

action questions, but not for object-property questions. As can be seen in Fig. 2, the results were as predicted: (1) there was no effect of form class on the answers concerning the first context (as subjects could answer correctly on the basis of the rich context); and (2) there was a strong effect of form class on answers for the second context. When the subjects read the nonce word as a noun, they stored with it properties of the objects in the context. When they read it as a verb, they stored with it properties of the action, notably the manner of action. Each of the four individual stories showed the same pattern, with the only difference being the absolute size of the percentages.

Separate chi-squared tests were performed for answers concerning the first and second contexts. As expected, for the first context there was no effect of form class. For the second context, there was a strong effect of form class ($\chi^2 = 31.25, p < 0.01$). Thus, we infer that when a new word was encountered in context, subjects used its form class to determine *which* aspects of the context they encoded as part of its meaning.

GENERAL DISCUSSION

The two studies presented here both bear on whether people use knowledge of the semantic regularities in their language to guide their hypotheses concerning word meanings. In Experiment 1, we tested a wide variety

of possible semantic regularities. We found that adults utilised a variety of types of knowledge in evaluating possible meanings for novel words encountered in context. There was evidence both for language-general constraints or biases (e.g. the taxonomic organisations of noun categories) and for language-specific semantic regularities (e.g. the structure of possible kinship terms, or the range of ancillary information that can be incorporated into a verb).

In the first study we asked people directly about the plausibility of possible word meanings. In Experiment 2, through the use of extended texts and questions which indirectly assessed word meanings, we indirectly assessed the word meanings that subjects derived from reading a new word in context. The results again indicated that subjects use knowledge of language-specific regularities in making hypotheses about the meanings of new words encountered in context. In this experiment, we believe we have closely approximated the normal process of learning word meanings from written context.

It is important to note that most of the constraints we have considered here are unlikely to be either innate or absolute (e.g. the word *contraband* probably constitutes a genuine exception to one of the regularities tested in Experiment 2). Rather, our point has been to establish the role of specifically linguistic knowledge in the formulation of hypotheses about the meanings of new words.

Knowledge about what constitutes a possible word meaning in one's language is seldom brought to conscious attention, so it is easy to underestimate its importance in vocabulary acquisition. However, the efficiency with which adults seem to learn word meanings from context, and the developmental evidence for knowledge of possible word meanings in children (Bowerman, 1988; Carey, 1978; Heilbeck & Markman, 1987) suggest that language learners must somehow acquire substantial knowledge about allowable word meanings. However, we are only beginning to gather direct evidence concerning this kind of knowledge. Research in developmental psychology on a small number of semantic constraints needs to be extended both to a larger set of constraints and to the use of these constraints in adult word learning. In this study we have tried to begin this extension. The results of our two experiments confirm that adults possess implicit knowledge of a wide variety of constraints of possible word meanings and that they apply this knowledge both in direct judgements of the plausibility of word meanings and in learning word meanings from context.

The success of this methodology suggests a number of further issues to investigate. An important developmental issue is the order in which constraints are acquired. We conjecture that universal constraints may be acquired before language-specific constraints. One example of this "universality predicts priority" principle (Bowerman & Gentner, in prep.) is

the finding discussed earlier, that children are slow to learn the allowable verb semantic patterns in their language (Gentner, 1982); such patterns (e.g. whether verbs of motion include manner or direction as part of their meaning) tend to be language-specific. Similarly, Bowerman (1988) has suggested that children may be relatively slow to learn such language-specific covert classes as the set of verbs that can take the *un-* prefix (e.g. *unload*, *untie*). The techniques explored in this research, particularly in Experiment 2, should be applicable to testing such developmental hypotheses.

A second important line of investigation that follows from this research is cross-linguistic comparison, including second-language learning. Language-specific semantic constraints may be important in predicting the translatability of different aspects of linguistic system. Further, an important source of difficulty for second-language learners may be the compatibility of the semantic constraints from the initial language with those of the second language. An indirect line of evidence for these cross-linguistic predictions is the observation that verbs are often slower to be acquired in a second language than are nouns (e.g. Dietrich, 1985). If, as Gentner (1981a; 1982) has suggested, verb meanings tend to vary more cross-linguistically than noun meanings, the relatively slow learning of verbs in a second language might reflect learners' difficulties with a new semantic system.

In summary, knowledge of the allowable semantic patterns of word meaning is an area of research that, although in its infancy, may have important implications for applications such as second-language learning, for theories of the acquisition of word meaning, and, most importantly, for our understanding of the nature of human linguistic knowledge.

Manuscript received July 1988

Revised manuscript accepted December 1989

REFERENCES

- Beck, I., McKeown, M., & McCaslin, E. (1983). All contexts are not created equal. *Elementary School Journal*, 83, 177-181.
- Bowerman, M. (1981). The child's expression of meaning: Expanding relationships among lexicon, syntax, and morphology. In H. Winitz (Ed.), *Proceedings of the New York Academy of Sciences Conference on Native Language and Foreign Language Acquisition*. New York: New York Academy of Sciences.
- Bowerman, M. (1982a). Reorganizational processes in lexical and syntactic development. In E. Wanner & L. R. Gleitman (Eds.), *Language acquisition: The state of the art*. Cambridge: Cambridge University Press.
- Bowerman, M. (1982b). Evaluating competing linguistic models with language acquisitional data: Implications of developmental errors with causative verbs. *Quaderni di Semantica*, 3, 5-66.

- Bowerman, M. (1988). Inducing the latent structure of language: The child's analysis of the input and the investigator's analysis of the child. In F. Kessel (Ed.), *The development of language and language researchers: Essays presented to Roger Brown*. Hillsdale, N. J.: Lawrence Erlbaum Associates Inc.
- Bowerman, M. & Gentner, D. (in preparation). A cross-linguistic study of the acquisition of relational terms: The case of Dutch and English *on*.
- Brown, R. (1957). Linguistic determinism and part of speech. *Journal of Abnormal and Social Psychology*, 55, 1-5.
- Carey, S. (1978). The child as word learner. In M. Halle, J. Bresnan, & G. A. Miller (Eds), *Linguistic theory and psychological reality*. Cambridge, Mass.: MIT Press.
- Carey, S. (1983). Constraints on word meaning. In T. Seiler & W. Wannenmacher. (Eds). *Concept development and the development of word meanings*. New York: Springer-Verlag.
- Casad, E. H. & Langacker, R. W. (1982). "Inside" and "outside" in Cora grammar. Unpublished paper.
- Clark, E. V. (1981). Lexical innovations: How children learn to create new words. In W. Deutsch (Ed.), *The child's construction of language*. London: Academic Press.
- Clark, E. V. (1982). The young word-maker: A case study of innovation in the child's lexicon. In E. Wanner & L. Gleitman (Eds), *Language acquisition: The state of the art*. Cambridge: Cambridge University Press.
- Clark, E. V. (1983). Meanings and concepts. In J. Flavell & E. Markman (Eds), *Handbook of child psychology, Vol. 3: Cognitive development*, pp. 787-840. New York: John Wiley.
- Clark, E. V. & Clark, H. H. (1979). When nouns surface as verbs. *Language*, 55, 767-811.
- Clark, E. V. & Hecht, B. F. (1982). Learning to coin agent and instrumental nouns. *Cognition*, 12, 1-24.
- Deighton, L. (1959). *Vocabulary development in the classroom*. New York: Bureau of Publications, Teachers College, Columbia University.
- Dickinson, D. (1984). First impressions: Children's knowledge of words gained from a single exposure. *Applied Psycholinguistics*, 5, 359-373.
- Dietrich, R. (1985). *Nouns and verbs in the learner's lexicon*. Paper presented at the Workshop on European Second Language Acquisition Research, The Linguistic Institute, Georgetown Heights.
- Fillmore, C. (1978). On the organization of semantic information in the lexicon. In D. Farkas, W. Jacobsen, & K. Todrys (Eds), *Papers from the parasession on the lexicon*. Chicago: Chicago Linguistic Society.
- Friedrich, P. (1970). Shape in grammar. *Language*, 46, 379-407.
- Gentner, D. (1975). Evidence for the psychological reality of semantic components: The verbs of possession. In D. A. Norman, D. E. Rumelhart, & the LNR Research Group, *Explorations in cognition*, pp. 211-246. San Francisco, Calif.: W. H. Freeman.
- Gentner, D. (1978). On relational meaning: The acquisition of verb meaning. *Child Development*, 49, 988-998.
- Gentner, D. (1981a). Some interesting differences between nouns and verbs. *Cognition and Brain Theory*, 4, 161-178.
- Gentner, D. (1981b). Verb semantic structures in memory for sentences: Evidence for componential representation. *Cognitive Psychology*, 13, 56-83.
- Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S. Kuczaj (Ed.), *Language development: Language, cognition and culture*. Hillsdale, N.J.: Lawrence Erlbaum Associates Inc.
- 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*. Cambridge: Cambridge University Press.
- Goldin-Meadow, S., Seligman, M., & Gelman, R. (1976). Language in the two-year-old. *Cognition*, 4, 189-202.

- Heibeck, T. & Markman, E. (1987). Word learning in children: An examination of fast mapping. *Child Development*, 58, 1021-1034.
- Hutchinson, J. (1985). *A constraint on children's implicit hypotheses about word meanings*. Paper presented at the meeting of the Society for Research in Child Development, Toronto, April.
- Jackendoff, R. (1975). A system of semantic primitives. In R. Schank & B. L. Nash-Webber (Eds), *Proceedings of the conference on theoretical issues in natural language processing*, pp. 24-29. Distributed by The Association for Computational Linguistics.
- Jackendoff, R. (1983). *Semantics and cognition*. Cambridge, Mass.: MIT Press.
- Jenkins, J. & Dixon, R. (1983). Vocabulary learning. *Contemporary Educational Psychology*, 8, 237-260.
- Katz, N., Baker, E., & MacNamara, J. (1974). What's in a name? A study of how children learn common and proper nouns. *Child Development*, 65, 469-473.
- Keil, F. (1979). *Semantic and conceptual development: An ontological perspective*. Cambridge, Mass.: Harvard University Press.
- Keil, F. (1981). Constraints on knowledge and cognitive development. *Psychological Review*, 88, 197-227.
- Keil, F. (1983). Semantic inferences and the acquisition of word meaning. In T. Seiler & W. Wannemacher (Eds), *Concept development and the development of word meaning*. New York: Springer-Verlag.
- Kintsch, W. (1986). Learning from text. *Cognition and Instruction*, 3, 87-108.
- Lakoff, G. (1987). *Women, fire, and dangerous things*. Chicago: University of Chicago Press.
- Landau, B. & Gleitman, L. (1985). *Language and experience*. Cambridge, Mass.: Harvard University Press.
- Langacker, R. (1982). Space grammar, analysability, and the English passive. *Language*, 58, 22-80.
- Langacker, R. (1986). *Foundations of cognitive grammar, Vol. 1: Theoretical prerequisites*. Stanford, Calif.: Stanford University Press.
- Maratsos, M. & Chalkley, M. (1981). The internal language of children's syntax. In K. E. Nelson (Ed.), *Children's language*, Vol. 2. New York: Gardner Press.
- Markman, E. (1987). How children constrain the possible meanings of words. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization*. Cambridge: Cambridge University Press.
- Markman, E. & Hutchinson, J. (1984). Children's sensitivity to constraints on word meaning: Taxonomic versus thematic relations. *Cognitive Psychology*, 16, 1-27.
- Miller, G. A. & Gildea, P. (1987). How children learn words. *Scientific American*, 257(3), 94-99.
- Miller, G. & Johnson-Laird, P. (1976). *Perception and language*. Cambridge, Mass.: Harvard University Press.
- Nagy, W. (1978). Some non-idiom larger-than-word units in the lexicon. In D. Farkas, W. Jacobsen, & K. Todrys (Eds), *Papers from the parasession on the lexicon*. Chicago: Chicago Linguistic Society.
- Nagy, W. & Anderson, R. C. (1984). How many words are there in printed school English? *Reading Research Quarterly*, 19, 304-330.
- Nagy, W. & Herman, P. (1987). Breadth and depth of vocabulary knowledge: Implications for acquisition and instruction. In M. McKeown & M. Curtis (Eds), *The nature of vocabulary acquisition*. Hillsdale, N.J.: Lawrence Erlbaum Associates Inc.
- Nagy, W., Herman, P., & Anderson, R. C. (1985). Learning words from context. *Reading Research Quarterly*, 20, 233-253.
- Nagy, W., Anderson, R. C., & Herman, P. (1987). Learning words from context during normal reading. *American Educational Research Journal*, 24, 237-270.

- Nelson, K. (1973). Structure and strategy in learning to talk. *Monographs for the Society for Research in Child Development*, 38 (Serial No. 149, Nos 1-2).
- Nelson, K. (1988). Constraints on word learning? *Cognitive Development*, 3, 221-246.
- Pinker, S. (1984). *Language learnability and language development*. Cambridge, Mass.: Harvard University Press.
- Plank, F. (1980). *Verbs and objects in semantic agreement: Minor differences between languages that might suggest a major one*. Unpublished manuscript. Hanover, West Germany.
- Quine, W. (1960). *Word and object*. Cambridge, Mass.: MIT Press.
- Sapir, E. (1944). *Language*. New York: Harcourt, Brace and World.
- Schatz, E. & Baldwin, R. S. (1986). Context clues are unreliable predictors of word meanings. *Reading Research Quarterly*, 21, 439-453.
- Slobin, D. (1973). Cognitive prerequisites for the development of grammar. In G. Ferguson & D. Slobin (Eds), *Studies of child language development*. New York: Holt, Rinehart and Winston.
- Soja, N., Carey, S. & Spelke, E. (1985, April). *Constraints on the meanings of words*. Paper presented at the meeting of the Society for Research in Child Development, Toronto.
- Sommers, F. (1963). Types and ontology. *Philosophical Review*, 72, 327-363.
- Talmy, L. (1972). *Semantic structures in English and Atsugewi*. Unpublished doctoral dissertation, University of California, Berkeley.
- Talmy, L. (1975). Semantics and syntax of motion. In J. Kimball (Ed.), *Syntax and semantics*, Vol. 4. London: Academic Press.
- Talmy, L. (1978). The relation of grammar to cognition. In D. Waltz (Ed.), *Proceedings of the second conference on theoretical issues in natural language processing*. New York: Association for Computing Machinery.
- Templin, M. (1957). Certain language skills in children: Their development and interrelationships. *The Institute of Child Welfare, Monograph Series No. 26*. Minneapolis: University of Minnesota.
- Werner, H. & Kaplan, E. (1952). The acquisition of word meanings: A developmental study. *Monographs of the Society for Research in Child Development*, 15 (Serial No. 51, No. 1).
- Wykes, T. & Johnson-Laird, P. N. (1977). How do children learn the meanings of verbs? *Nature*, 268, 326-327.

APPENDIX 1

List of Constraints Investigated

For each hypothesised constraint on English word meanings considered in this experiment, there is a brief description of the constraint, followed by the two contextually appropriate phrases from the item(s) testing that constraint. The phrase violating the hypothesised constraint is marked with an asterisk, indicating that it should not constitute a well-formed word meaning in English. We also note whether we would expect this constraint to reflect a universal constraint on word meanings, or a possible restriction on word meanings in English.

Constraints on Nouns

1. *Taxonomic constraint on nouns* (universal). Noun meanings tend to reflect taxonomic relationships rather than thematic relationships. For example, there is more likely to be a single word meaning "needles and pins" than a word meaning "needles and thread".

*native foods and cooking utensils

vs small nocturnal mammals with scent glands

*things like needles, thread and scissors

vs things like pieces of silk, bits of flowered cotton, and scraps of corduroy

2. *Durative constraint on nouns* (universal). Nouns are usually differentiated in terms of relatively long-term properties rather than relatively short-term properties. For example, there is more likely to be a term for "red bird" than for "flying bird". We expect this constraint to reflect a universal tendency. However, it clearly applies to morphologically simple nouns, and not to deverbal nouns such as *batter* (i.e. the person currently at bat).

*angry baboon

vs poisonous snake

3. *Transparency of agency* (universal). Nouns that have an agentive meaning, i.e. "person who does X", usually are fully transparent in English, e.g. *letter carrier* or *mailman*. Those that are not fully transparent, e.g. *doctor* or *dentist*, still have an overt agentive suffix. We expect this to be a universal tendency, but with exceptions, e.g. *architect* and *secretary* in English.

*furniture carriers

vs groups of soldiers

4. *No specification of linearity in kinship terms* (English-specific). English does not have kinship terms which specify side of the family (father's or mother's). There are clearly other languages in which this constraint does not apply.

*mother's parents

vs the people next door

5. *No specification of age in sibling terms* (English-specific). English does not have terms which distinguish siblings by relative age. There are languages which have such terms (e.g. a word meaning "younger brother") which are not possible in English.

*younger brothers and sisters

vs nieces and nephews

Constraints on Verbs

6. *No incorporation of patients in verbs* (English-specific). English verbs do not include specifications of particular properties of patients (i.e. typically, direct objects). For example, English does not have, and is not likely to acquire, a word meaning "to wave a long rigid object", or a word meaning "to eat crackers". Languages appear to differ in the extent to which they allow the incorporation of properties of objects into verbs. (Note that English verbs can incorporate the meanings of *instruments*, especially when they are overtly marked in the morphology, as in the verb *to hammer*: see Clark and Clark, 1979.)

*to eat snails

vs to eat with chopsticks

*to spill some dirt

vs to roll end over end

7. *No specification of time in verbs* (English-specific). English verbs do not incorporate information about clock or calendar time (i.e. time of day or season of the year, as opposed to tense) that the event referred to occurs:

*to spend the afternoon
vs to walk aimlessly

*to wake up early
vs to wake up quickly

*to arrive exactly on the hour
vs to tip generously

8. *No cessation-of-X constraint on verbs* (weakly universal). English verbs do not incorporate the meaning "cessation" along with a specific activity; there are verbs like *stop* and *cease*, but not verbs with a meaning like (for example) "to stop eating".

*to stop sewing
vs to laugh nervously

9. *No continuation-of-X-constraint on verbs* (English-specific). English verbs do not usually incorporate a component of iteration or continuation of an action. That is, there is not likely to be a verb distinct from *dance* that means "to keep on dancing". There are at least some apparent exceptions in English, e.g. *beat* vs *hit*.

*keep on inspecting
vs quickly grab

10. *No combination of actions constraint on verbs* (universal). Verb meanings tend not to include two distinct actions. This is presumably a universal tendency, although it is not clear that this constraint can avoid circularity, given the difficulty of finding any non-linguistic criteria for what constitute "distinct actions".

*to soak and bend
vs to get completely ready

*to run around singing
vs to run around frantically

11. *Constraints on the incorporation of direction into verbs of motion* (English-specific). English motion verbs only infrequently incorporate direction into their meaning, and there are tight restrictions on what type of derivational meanings can be incorporated. Directional information also represented independently by prepositions can be incorporated, as in *enter*, *exit*, *transverse*, *penetrate*, *leave*, *arrive*, as can vertical direction, as in *rise*, *descend*. However, other types of directional information, e.g. cardinal points of the compass, cannot be incorporated into verb meanings.

*to go north
vs disappear quickly

12. *No incorporation of direction in verbs of speaking* (English-specific). English verbs of speaking do not incorporate information about direction.

*to shout over (a barrier)
vs to jump over (a barrier)

13. *No incorporation of speed in verbs of thought* (English-specific). English verbs referring to mental processes tend not to include manner specifications concerning speed.

*to make decisions quickly
vs move very quickly

14. *No incorporation of mood in verbs of writing* (English-specific). English verbs of writing, unlike verbs of speaking, do not incorporate specifications of mood.

*to write angrily
vs to shout defiantly

Constraints Across Part-of-speech Categories

15. *Transparency of meanings involving proper names* (weakly universal). Proper names are not included in a word's meaning unless they are present in its form. For example, *Americanise* is a legitimate English word, and *Frenchify* is a possible English word, but there could not be a morphologically opaque word, e.g. *rarificate*, meaning 'to translate into Polish'.

*drive around in a Buick
vs take a leisurely tour.

APPENDIX 2

Sample Passage from Experiment 2: "The Darsts of Zurdine"

Noun Version

The planet Zurdine has the unique feature that its plants grow very fast and its animals move very slowly. In fact, many of its plant forms grow more quickly than any of its animal forms are able to move. This fact has led to some very interesting adaptations, especially among the Darsts, a humanoid species inhabiting the plains of Zurdine. The Darsts, although surrounded by fast-growing plants, are intelligent enough to capitalise upon their environment.

Fortunately for the Darsts, the planet is extremely dry, so by controlling the available water they can manage plant growth to some extent. Their way of life depends heavily on the use of the seeds of the Yrang plant, which begin to sprout between one and two seconds after they have been exposed to any moisture, and grow at speeds of up to 15 feet a second. For defence, a hunting party can erect a temporary enclosure simply by scattering Yrang seeds and spitting around the perimeter. This is often a necessary defence as there are several types of dangerous animals on Zurdine. Some of the worst are large six-legged predators, which, although slow-moving, are relentless, and the roving herds of poisonous razor-spined herbivores.

There are also a number of hunting techniques involving Yrang seeds. The most common is simply to throw moistened Yrang seeds around the quarry, so that the animal is quickly enveloped in a net of vines.

I observed other techniques as well while on a hunting expedition with the Darsts. Our goal that day was to bring in a particular species of smaller rodent-like animals which are prized as a special delicacy by the Darsts. But so far we had seen nothing but larger predators, which we had been careful to avoid. Finally, after several hours, we caught sight of a small animal moving to our right. While the rest of the party stood motionless, the head of the hunting party pulled out a hollow reedlike peashooter, and moistened some Yrang seeds with saliva. She slipped them into the peashooter and with accuracy that amazed me, she ensnared the unsuspecting *sapryn*. The entire hunting party was soon ambling in that direction to put their prey into a more permanent net.

Although I was untrained in Darst hunting techniques, I was able to be of help at a moment of real danger. One of the younger Darsts, who move even more slowly than their languid elders, was somewhat behind the group, and not even making much effort to stay close.

Without warning, the young Darst found himself face to face with a *sapryn*. The youngster tried to ensnare it, but in his inexperience he dropped several moistened seeds. By the time I got to him, a vine was wrapped chokingly tight around his neck. I pulled him free, and carried him back to the group.

Verb Version

The planet Zurdine has the unique feature that its plants grow very fast and its animals move very slowly. In fact, many of its plant forms grow more quickly than any of its animal forms are able to move. This fact has led to some very interesting adaptations, especially among the Darsts, a humanoid species inhabiting the plains of Zurdine. The Darsts, although surrounded by fast-growing plants, are intelligent enough to capitalise upon their environment.

Fortunately for the Darsts, the planet is extremely dry, so by controlling the available water they can manage plant growth to some extent. Their way of life depends heavily on the use of the seeds of the Yrang plant, which begin to sprout between one and two seconds after they have been exposed to any moisture, and grow at speeds of up to 15 feet a second. For defence, a hunting party can erect a temporary enclosure simply by scattering Yrang seeds and spitting around the perimeter. This is often a necessary defence as there are several types of dangerous animals on Zurdine. Some of the worst are large six-legged predators, which, although slow-moving, are relentless, and the roving herds of poisonous razor-spined herbivores.

There are also a number of hunting techniques involving Yrang seeds. The most common is simply to throw moistened Yrang seeds around the quarry, so that the animal is quickly enveloped in a net of vines.

I observed other techniques as well while on a hunting expedition with the Darsts. Our goal that day was to bring in a particular species of smaller rodent-like animals which are prized as a special delicacy by the Darsts. But so far we had seen nothing but larger predators, which we had been careful to avoid. Finally, after several hours, we caught sight of a small animal moving to our right. While the rest of the party stood motionless, the head of the hunting party pulled out a hollow reedlike peashooter, and moistened some Yrang seeds with saliva. She slipped them into the peashooter and with accuracy that amazed me, she *sapryned* the unsuspecting animal. The entire hunting party was soon ambling in that direction to put their prey into a more permanent net.

Although I was untrained in Darst hunting techniques, I was able to be of help at a moment of real danger. One of the younger Darsts, who move even more slowly than their languid elders, was somewhat behind the group, and not even making much effort to stay close. Without warning, the young Darst found himself face to face with an animal. The youngster tried to *sapryn* it, but in his inexperience he dropped several moistened seeds. By the time I got to him, a vine was wrapped chokingly tight around his neck. I pulled him free, and carried him back to the group.