Word Learning

J Snedeker, Harvard University, Cambridge, MA, USA

Word learning is often considered the simplest and least controversial aspect of language development. Although theorists fiercely debate the ontogenetic and phylogenetic origins of grammar, everyone agrees that words must be learned by observing the contexts in which they are used. No other theory can explain how English-speaking children come to use 'shoe' to label footwear, whereas young French speakers use the same sequence of sounds to label cabbage. However, this self-evident truth masks a host of questions about how learning occurs and the knowledge that children bring to the problem.

Defining the Problem

To understand word learning, we need to know both what a word is and which aspects of words are learned. Surprisingly, neither question has a clear answer. There are two difficulties in defining what a word is.

The first is determining what information is included as part of a word. By all definitions, each word (or lexical entry) consists of a symbol that is paired with a concept. In the case of spoken languages, these symbols consist of phonological representations. For example, my lexical entry for 'cat' includes the phoneme sequence /kæt/ which is linked to some conceptual representation of catness. Literate individuals also store the written form of words they are familiar with. In addition, language users have other knowledge of words that could be a part of the lexical entry. Much of this knowledge is related to the grammatical environments in which a word appears. For example, our knowledge of English allows us to recognize that "I have a cat" is a grammatical sentence, whereas "I like to cat" is not. In dictionaries, we capture these facts by assigning words to grammatical categories which are listed with each definition (e.g., 'cat' is a noun and not a verb). Many have argued that grammatical categories are also stored in the mental lexicon. However, our knowledge of the relation between lexical items and grammatical structure extends beyond syntactic categories. For example, we know that a verb such as 'roll' can appear in a transitive sentence ("I rolled the ball"), whereas a verb such as 'fall' cannot ("I fell the ball"). This has led many to propose that individual lexical entries contain detailed grammatical information about the contexts of word use.

Others have argued that the concepts with which words are associated allow us to derive these structural facts, eliminating the need for their storage.

The second definitional difficulty is determining how to individuate words: Which chunks of linguistic material are entered in our mental lexicons? How large are these chunks, and where does one entry end and the next begin? Clearly, words such as 'cat' that consist of a single morpheme (meaningful unit) will have to be stored because they cannot be derived in any way. However, we also store morphemes that are not whole words (-ness or -ed), some words that are made of multiple morphemes (e.g., 'walked' or 'happiness'), and idioms that consist of multiple words ('kick the bucket'). The fact that many phonological forms have several meanings creates further uncertainty about what constitutes a lexical entry. In the case of homonyms, these meanings are not related and clearly must be stored separately (e.g., river bank vs. savings bank). However, in the case of polysemes, the different meanings are related (e.g., line up vs. telephone line), and it is unclear whether the alternatives are two separate words, two ways of thinking about a single lexical entry, or two subentries of a single lexical item. Problems of this kind have led some theorists to reject the metaphor of the lexicon as the list of entries and to argue instead that the lexicon is a generative system for linking forms to meanings.

Once we have defined a word, we can ask what it means to learn one. Which aspects of a lexical entry are learned and which can be derived? Which parts of the entry precede word learning and which are the result of it? All theorists agree that children must learn the mapping between the phonological form and the concept. Although there are some minor islands of systematicity (e.g., in English words beginning with *sn-* are related to the nose), these mappings are largely arbitrary and vary across languages; thus, they must be learned. There is more controversy about the role of learning in our ability to represent the phonological form of the word or to represent the concept with which it is associated.

Adults and older children clearly represent speech in terms of stable phonological categories and thus simply have to learn the sequence of phonemes that is linked to a particular concept. The nature of infants' phonological representations has been more controversial. Because infants sometimes fail to distinguish between two novel words that are similar but phonologically distinct (e.g., 'bim' and 'dim'), it has been suggested that infants have less precise, gestalt representations of words and only develop phonological

representations as their lexicons grow. However, infants can succeed in learning phonologically similar words in more supportive contexts and have more difficulty comprehending a known word when a single phoneme has been altered. One possible interpretation is that infants represent word forms in the same way as older children but have difficulty attending to and encoding these forms when the task is challenging.

Many developmental psychologists have argued that word learning involves concept acquisition. Children do not merely map words to existing concepts; instead, they use the recurring word form as a cue to form a new category. This theory is compatible with the Whorfian hypothesis because the concepts that the learner acquires depend on the concepts that are lexicalized in the learner's language. This theory could potentially explain why children learn concrete words early and more abstract words later, or why labels for objects generally appear before predicates. However, theories of this kind face two challenges. First, it is unclear how a child could learn what a word means if he or she did not already have access to the concept that the word encodes. To recognize that a word reliably co-occurs with a particularly category of entities, one must be able to represent those entities as a class. In other words, you cannot not learn that /kæt/ means cat unless you already have a representational arsenal that allows you to distinguish cats from non-cats. Second, the course of lexical development in cognitively advanced learners suggests that conceptual limitations play little role in early word learning. Children who are internationally adopted as preschoolers must learn a new language in context similar to that of infant language learners. If the course of infant word learning is shaped by the conceptual limitations of babies, then we would expect that older children, free from such limitations, would be able to learn more verbs and abstract words during the initial phases of lexical development. However, internationally adopted preschoolers acquire words in approximately the same order as infant learners, suggesting that early vocabulary development is largely independent of cognitive development.

Word Learning across Development

Children learning both signed and spoken languages generally produce their first word at approximately 12 months of age. By adulthood, the typical person knows approximately 60 000 words. Although this averages out to approximately 9 words a day, the actual pace of word learning varies across development. As Figure 1 illustrates, the initial pace of word learning is quite slow, with the average child learning approximately 2 words a week between 12 and


Figure 1 Vocabulary growth throughout development. The vocabulary estimates for 12- to 30-month-olds are from Fenson L, Dale PS, Reznick JS, Bates E, Thal DJ, and Pethick SJ (1994) Variability in early communicative development. *Monographs of the Society for Research in Child Development* 59: 1–173. The estimates for 6- to 10-year-olds are from Anglin J (1993) Vocabulary development: A morphological analysis. *Monographs of the Society for Research in Child Development* 238: 1–166. The estimate for adults is based on the figure cited by Aitchinson J (1994) *Words in the Mind: An Introduction to the Mental Lexicon*, 2nd edn. Oxford, UK: Blackwell.

16 months of age. Until recently, many observers believed that there was a sudden acceleration in word learning at approximately 18 months of age (often called the vocabulary spurt). However, closer examination has revealed that the pace of word learning increases gradually throughout the toddler and preschool years. In literate societies, this steady acceleration continues into the school years as children encounter new words while reading. The pace of word learning begins to decelerate sometime between 8 and 18 years, presumably because older children encounter fewer unknown words. There is no evidence for a critical period in word learning: Adults do as well or better than children in most experimental word learning tasks and we can readily learn new words throughout our lives.

As children's vocabularies grow, they also change. The first words that children learn are typically labels for people ('mommy' or 'daddy'), animal names ('kitty'), social words ('hi' or 'uh-oh'), or utterances used in common routines ('peek-a-boo'). Although verbs, prepositions, and other relational words

appear in early vocabularies, they may have more limited meanings for infants than they do for adults and older children. For example, many 1-year-olds say 'more' to request the recurrence of an event or object, but there is no evidence that children of this age understand that 'more' is a term that quantifies the amount in one set relative to another.

The early vocabularies of children learning English are dominated by nouns, most of which label people, animals, or objects. Even though adults speak to young children in full sentences, complete with verbs and grammatical words, these elements are massively underrepresented in children's early vocabularies. The proportion of nouns in the child's lexicon varies from language to language depending on whether subjects and objects are mandatory and whether the verb appears in a perceptually salient position (e.g., at the beginning or end of the utterance). Nevertheless, in every language that has been systematically studied, nouns are overrepresented in children's vocabularies relative to their rate of occurrence in the input, whereas verbs are underrepresented. This suggests that nouns, particularly those that label objects or people, are easier for children to learn than verbs or other relational words. Some theorists have attributed this to the greater conceptual complexity of verbs, whereas others attribute it to the nature of the information needed to identify the meanings of nouns and verbs.

There is considerable variability in how rapidly children learn words. At 16 months of age, the median vocabulary size for middle-class children in the US, is approximately 40 words. However, 10% of children produce no words, whereas another 10% produce more than 150 words. Early in development, vocabulary size is a better predictor of vocabulary composition than age. As the child's vocabulary grows from 50 to 200 words, the proportion of these words that are nouns increases. Verbs and adjectives begin to appear in greater numbers between 100 and 400 words. Grammatical words (such as articles, pronouns, prepositions, and auxiliaries) increase in frequency at approximately 400 words. Between 16 and 30 months, the size of a child's productive vocabulary is tightly correlated with the grammatical complexity of the child's speech. Initially, children primarily produce one-word utterances. When their productive vocabulary reaches 50–200 words, they begin combining words into short phrases. From 300 to 500 words, grammatical morphemes appear and the child's utterances increase in length and complexity. These relations are preserved in internationally adopted preschoolers, suggesting that the correlation is not a side effect of global maturational or cognitive changes. In bilingual children, these relations hold within a language but not between languages, suggesting that lexical development facilitates grammatical development or vice versa. During the school years, lexical development and reading are closely linked. A child's vocabulary when he or she enters school is a strong predictor of later reading achievement, and subsequent vocabulary development is correlated with the amount that the child reads.

Word Learning as Induction

How do children learn the meanings of words? Most people who have thought about the problem long enough have come up with essentially the same solution: Word learning is a form of induction. Learners generate hypotheses based on the situation in which the new word occurs. As the learner observes new instances of the word, hypotheses are eliminated or strengthened, allowing him or her to close in on the correct meaning.

Although this mechanism clearly plays a role in word learning, it cannot be the entire story. First, even very young children can learn some words after hearing them used in just one context. Second, many words, particularly verbs and other relational terms, are often used in the absence of the event being labeled (e.g., when parents tell children to "go to sleep," both parties are typically awake).

Finally, the account of word learning given previously is subject to the mid-century critiques of empiricism voiced by philosophers such as W. V. O. Quine and Nelson Goodman. Learning simply cannot be unconstrained induction because any finite set of observations is consistent with an infinite set of hypotheses. To take Quine's example, the set of observations that would allow one to learn that a word ('gavagai') means rabbit is also compatible with the hypothesis that it means undetached rabbit parts (or temporal rabbit slices). Thus, a full account of lexical development requires more than merely stating that word learning is induction. In the past 25 years, developmental psycholinguists have begun to flesh out this story by identifying three ways in which children and adults can tame the induction problem and learn the meanings of words.

Constraints and Biases on Hypothesized Meanings

The first way to tame induction is by limiting the range of meanings that are considered as possible hypotheses. Many theoretically possible meanings may be ruled out simply because children cannot or do not think of them – for word learning or any other purpose. Quine's example of undetached rabbit parts

is presumably ruled out by cognitive constraints of this kind. However, such constraints cannot explain how the learner rules out more plausible alternatives such as white, fluffy, hopping, tail, animal, or Flopsie. All these hypotheses are plausible ones, meanings which the learner could eventually link to some other phonological form. Children, however, are biased learners who privilege some of these hypotheses over others.

The nature of these biases has been explored with an experimental paradigm called the word extension task. In this task, children are shown a novel object (e.g., a blue kidney-shaped piece of plaster) which is labeled with a novel word ("This is my dax"). Then they are given a small set of test objects (e.g., a blue circular piece of plaster and a red kidney-shaped piece of wood) and asked to find another example of the target word ("Can you give me the dax?"). The child's response allows the experimenter to infer how the child is interpreting the novel word.

Through tasks such as this one, researchers have discovered that children (and adults) have a strong bias to assume that a new word refers to a whole object rather than one of its parts, its properties, or the relations that it is involved in. When children have mapped a word to an object, they will extend it to other objects of the same kind rather than to other objects that are involved in the same event (e.g., extending from a birthday cake to other cakes rather than to birthday presents and birthday candles). In the case of novel artifacts, they typically extend the word to objects of the same shape rather than to objects of the same material or same color. However, when a word is applied to a novel animal, children are more conservative, preferring referents that have the same color or texture and the same shape.

These biases should aid the child in learning count nouns (such as 'cat,' 'book,' or 'car'), but they might hinder the acquisition of adjectives, verbs, relational nouns, and labels for substances and parts of objects. For example, if little Johnny were to hear the word 'silver' applied to Mommy's car, his bias to map words to whole objects could lead him to incorrectly conclude that 'silver' means car. One additional constraint, often called mutual exclusivity, may help learners overcome this bias. Young children are reluctant to map a second label to an object for which they already have a label. Thus, if little Johnny has already learned that cars are called 'car,' he is likely to reject the hypothesis that 'silver' means car and consider other possibilities.

The existence of these biases is uncontroversial, but their origins are not. Current research in this area focuses on two issues. The first is domain specificity. Are these biases unique to word learning or do they reflect more general properties of cognition? Perhaps

the bias to link words to whole objects reflects the centrality of objects in early cognition rather than any preconceptions about the nature of words. Similarly, mutual exclusivity in word learning could result from a more general principle of communication which leads learners to assume that speakers will use known forms of reference whenever possible. The second issue is whether these biases are available at the onset of word learning. Some theorists suggest that these biases are learned as the child acquires words and notices the properties that are typically used in word extension. For example, children may learn their first artifact labels by trial and error. After learning many of these words, a child may notice that these labels are generally extended by shape and develop a bias to extend new words in this way. Several pieces of evidence support this account. The shape bias for novel artifacts increases as vocabulary size grows, and toddlers who are taught categories that are extended on the basis of shape develop this bias more quickly and are more successful at learning count nouns.

Social Cues to Reference

Children can also tame the induction problem by using their implicit understanding of social interactions to make inferences about the communicative intentions of the speaker. These inferences can help the child identify the object or event to which the speaker is referring, thus simplifying the problem of word learning. The best studied social cue to word meaning is the speaker's direction of gaze. When speakers are talking about objects that are visible, they often look at that object at roughly the same time as they are mentioning it. By approximately 18 months of age, children will map a word to the object that the speaker was looking at, even when they themselves were examining a different object when the word was spoken. Infants can also identify the referent of a word when the speaker points at the object or moves it as it is labeled.

Other abilities appear to require a more complex representation of the interaction. When a speaker labels an action before performing it, young children prefer to link the label to an intentional action rather than an accidental one. When a speaker uses a novel label, children assume that she is referring to an object that was introduced while she was gone rather than an object that she saw before.

Research on the use of social cues has explored two questions. First, what role do these cues play in word learning? Are they necessary or merely facilitory? Several lines of evidence demonstrate that children can learn words in the absence of any single social cue. For example, infants can learn the mapping

between an object and a word from a videotape in which the person who is producing the word is never visible. The ability of blind children to acquire words at approximately the same age as sighted children provides further evidence that visual access to social cues (such as pointing or gaze direction) is not necessary for successful word learning. Of course, it is possible that these learning contexts contain other cues that allow the child to infer the referential intentions of the speaker. A second and related question is whether infants' use of social cues reflects an understanding of the mental states of the speaker rather than simple associations or low-level attentional processes. Do infants make use of eye gaze because they know that it reflects the speaker's knowledge and communicative intentions, or have they merely learned that gaze is a reliable predictor of which object a word will be associated with?

Sentential Contexts as Cues to Word Meanings

The final way in which children tame the problem of word learning is by using the sentential context in which a word occurs to narrow down the set of possible hypotheses. Sentential contexts provide three kinds of information. First, they identify the syntactic category of a novel word. Because syntactic categories are systematically (but imperfectly) linked to semantic categories, they provide information about the kind of meaning that the word is likely to have. For example, 1-year-olds who are shown a doll and told that 'this is a zav' interpret the novel word as a common noun and extend it to other similar dolls. In contrast, those who hear 'this is Zav' interpret it as a proper noun and refuse to extend it. Similarly, by approximately 2½ years of age, young children have learned to extend mass nouns ('this is some dax') to entities that are made of the same substance and to extend adjectives ('this is a daxy one') to objects that share a common property.

Second, sentential contexts provide information about the number and kinds of arguments that a predicate can take which constrain the kind of meaning that the predicate has. For example, verbs of self-generated motion (e.g., 'dance') often appear with just one argument (the subject); verbs of contact (e.g., 'hit') or caused motion (e.g. 'push') usually appear with a subject and a direct object; and verbs of transfer (e.g., 'send') appear with three arguments (subject, direct object, and indirect object). Young children use these connections between syntax and semantics to learn the meanings of unknown words. For example, 2-year-olds who hear a transitive utterance such as "She is blicking her around" expect the verb to have a meaning such as 'push,' whereas those

who hear an intransitive utterance such as "They are blicking around" expect the verb to have a meaning such as 'dance.'

Finally, sentential contexts contain known words which provide information about the other entities or relations that are involved in the event under discussion. For example, if children hear an utterance such as 'John is eating a dax,' they can infer that a dax is probably a kind of food, and they can direct their attention to John to ascertain precisely what kind of food it might be. Although cues of this kind tightly constrain word learning, their availability depends on the child having already learned some of the other content words in the utterance.

Current research on the role of sentential contexts centers on two issues. First, when do children begin to use each of these cues? Some cues have an influence on infants as young as 13 months of age (e.g., use of word in count noun syntax), but other cues appear to have little effect even in preschoolers. For example, in several studies 2- and 3-year-olds have failed to use verbal morphology ("Look gorping!") to map the word to the action (e.g., waving) that is being performed. The second active area of research explores how these cues are acquired. One possibility is that by learning words from particular categories, children discover a direct association between specific grammatical contexts and particular aspects of word meaning. The second possibility is that children's word learning is guided by universal linkages between syntax and semantics. Children must learn the grammatical morphemes that mark each syntactic category (e.g., that 'a dax' is a count noun in English and 'une dax' is a count noun in French) but the connection between particular syntactic categories and patterns of word extension comes for free. These theories make different predictions about when these cues should be learned and whether the semantic relations in question are likely to be linked to observable properties (e.g., count nouns pick out shape-based categories) or more abstract properties that are tightly linked to the syntax (e.g., count nouns pick out individuated entities).

Final Words

During the past 25 years, researchers have demonstrated that young children make use of all three of the previously discussed strategies to learn words. The challenge in the future will be to understand how these strategies work together throughout development. It may be useful to think of constraints, social cues, and structural contexts as complementary mechanisms rather than competing hypotheses. The three strategies account for different aspects of word learning. For example, social cues allow children to

identify the referent of the phrase, whereas constraints on hypotheses and structural cues provide guidance about how this referent is being construed.

The order in which these abilities emerge during development may explain why early vocabularies are dominated by nouns but contain few verbs, and why children's mastery of verbs coincides with the emergence of grammar. Many social cues and constraints may be available at the onset of word learning. These abilities do not necessarily depend on prior linguistic knowledge, and their use extends well beyond the linguistic domain (e.g., knowledge of object kinds constrains reasoning as well as word learning, and eye gaze provides information about a person's knowledge outside of the linguistic domain). These abilities would primarily facilitate the acquisition of nouns: Most word learning constraints and biases apply to object categories and most social cues function to pick out objects in space. In contrast, the use of structural contexts necessarily requires some grammatical knowledge and thus can only be used after language acquisition is underway. This information source is particularly useful for learning verbs and other predicates (which occur in richer structural contexts). Thus, efficient verb learning may emerge only when children have acquired enough grammatical and lexical knowledge to make use of sentential contexts during word learning.

See also: Bilingualism; Language Development; Language: Cortical Processes; Lexical Impairments Following Brain Injury; Second Language Acquisition; Word Production; Word Recognition.

Further Reading

Baldwin DA (1991) Infants' contribution to the achievement of joint reference. *Child Development* 62: 875–890.

Bates E, Dale PS, and Thal D (1995) Individual differences and their implications for theories of language development. In: Fletcher P and MacWhinney B (eds.) *The Handbook of Child Language*, pp. 96–151. Oxford, UK: Blackwell.

Bates E and Goodman JC (1997) On the inseparability of grammar and the lexicon: Evidence from acquisition, aphasia and real-time processing. *Language and Cognitive Processes* 12(5/6): 507–586.

Bloom P (2000) *How Children Learn the Meanings of Words.* Cambridge, MA: MIT Press.

Caselli MC, Casadio P, and Bates E (1999) A comparison of the transition from first words to grammar in English and Italian. *Journal of Child Language* 26: 69–111.

Gentner D and Boroditsky L (2001) Individuation, relativity and early word learning. In: Bowerman M and Levinson S (eds.) *Language Acquisition and Conceptual Development.* Cambridge, UK: Cambridge University Press.

Gillette J, Gleitman H, Gleitman L, and Lederer A (1999) Human simulation of vocabulary learning. *Cognition* 73: 135–151.

Gleitman LR (1990) The structural sources of word meaning. *Language Acquisition* 1: 3–55.

Hall DG and Waxman SR (2004) *Weaving a Lexicon.* Cambridge, MA: MIT Press.

Markman EM (1989) *Categorization and Naming in Children.* Cambridge, MA: MIT Press.

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

Smith LB, Jones SS, Landau B, Gershkoff-Stowe L, and Samuelson L (2002) Object name learning provides on-the-job training for attention. *Psychological Science* 13: 13–19.

Snedeker J, Geren J, and Shafto C (2007) Starting over: International adoption as a natural experiment in language development. *Psychological Science* 18: 79–87.

Waxman SR (1999) Specifying the scope of 13-month-olds' expectations for novel words. *Cognition* 70: B35–B50.