Cognition 117 (2010) 17–37

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

Shifting senses in lexical semantic development

Hugh Rabagliati [a,*], Gary F. Marcus [a], Liina Pylkkänen [a,b]

[a] Department of Psychology, New York University, 6 Washington Place, New York, NY 10003, United States
[b] Department of Linguistics, New York University, 10 Washington Place, New York, NY 10009, United States

ARTICLE INFO

Article history:
Received 26 June 2009
Revised 3 June 2010
Accepted 14 June 2010

Keywords:
Lexical development
Lexical semantics
Word learning
Language acquisition

ABSTRACT

Most words are associated with multiple senses. A DVD can be round (when describing a disc), and a DVD can be an hour long (when describing a movie), and in each case DVD means something different. The possible senses of a word are often predictable, and also constrained, as words cannot take just any meaning: for example, although a movie can be an hour long, it cannot sensibly be described as round (unlike a DVD). Learning the scope and limits of word meaning is vital for the comprehension of natural language, but poses a potentially difficult learnability problem for children. By testing what senses children are willing to assign to a variety of words, we demonstrate that, in comprehension, the problem is solved using a productive learning strategy. Children are perfectly capable of assigning different senses to a word; indeed they are essentially adult-like at assigning licensed meanings. But difficulties arise in determining which senses are assignable: children systematically overestimate the possible senses of a word, allowing meanings that adults rule unlicensed (e.g., taking *round movie* to refer to a disc). By contrast, this strategy does not extend to production, in which children use licensed, but not unlicensed, senses. Children's productive comprehension strategy suggests an early emerging facility for using context in sense resolution (a difficult task for natural language processing algorithms), but leaves an intriguing question as to the mechanisms children use to learn a restricted, adult-like set of senses.

1. Introduction

Human language is filled with the ambiguous and non-literal. When the witches of Macbeth urge that their *fire burn and cauldron bubble*, they do not mean for the cauldron itself to melt and boil, but the magic potion inside. Similarly, if I *order some Beethoven from the music store*, I have not ordered a lump of the composer, but rather some of his works; if I *find the CD to be moving*, it is the composer's works that cause emotion, not the plastic CD itself. In each example, the surface meaning of the sentence seems implausible but by shifting the meaning of a critical constituent we can derive a reasonable interpretation.

However, the elasticity of meaning only stretches certain ways. Although *Beethoven* can refer to the composer's music, his music cannot refer to him; it is nonsensical to say that *the 8th symphony was deaf*. Similarly, *the CD* can refer to the composer's work, but not vice versa (e.g., *the 8th symphony was shiny*).

The child learning a language has to figure out these ground rules, a task that is far from trivial. Computer scientists have spent forty years failing to create a computer program that can adequately determine the similar but different meanings associated with words such as *Beethoven* or *CD*, which linguists call senses (for an overview see Miller, 1999). Yet by adulthood our ability to resolve a word's sense is extremely accurate. How, then, do children learn the ways a word's sense can change?

* Corresponding author. Address: Department of Psychology, New York University, 6 Washington Place, 8th Floor, New York, NY 10003, United States. Tel.: +1 212 998 3870.
E-mail address: hugh@nyu.edu (H. Rabagliati).

0010-0277/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.cognition.2010.06.007

1.1. Shifts and senses

Although children eventually attain a remarkable degree of mastery over the scope and limits of word meaning, relatively little is known about how and when they manage to do so. We know a great deal about how children use logical principles, theory of mind, syntax and other factors to determine the basic referent of a word when heard for the first time, but little about how a child's understanding of a word extends beyond those first encounters. Carey and Bartlett (1978) argued that any word requires a long period of slow discovery before a child finally determines its exact meaning, a point that has been echoed by Murphy (2001). Presumably, learning the many ways a word's meaning changes is similarly difficult.

What sorts of representations do children have to acquire? One possibility is that the lexicon simply lists a set of word forms paired with their meanings. Entries for words such as *DVD* would contain the form alongside a sense referring to a shiny disc and a sense referring to the movie stored on that disc. This theory is attractively simple and to some extent may even be true: there is no reason why frequently used senses of a word could not be stored together. But it has difficulty accounting for any sort of creative word use. For example, an animal's name can also refer to the food produced from it (compare *noisy chicken* and *tasty chicken*), and this sort of template generalizes across words with similar meanings. A reader offered a steaming saucer of *Sasquatch*[1] would find the food sense entirely transparent yet also entirely novel. If they had lacked productive means for changing a word's sense then the only available meaning would have been the often-encountered animal sense, and not the novel food one (for discussion see Murphy, 2007).

Mastering this type of creativity requires the child to learn a number of productive shifts[2] that can generate senses based on certain aspects of a word's meaning. As shown below, we can use the name of an object to refer to its abstract contents (1), take a container's name to stand for its contents (2), or interpret a physical object as taking part in some unspecified event (3).

(1) The DVD was an hour long.
= *The movie on the DVD was an hour long.*
(2) The pot was stirred.
= *The content of the pot was stirred.*
(3) The boy began the book.
= *The boy began reading/writing the book.*

Under most theories, productive shifts in meaning are the result of rules applying over coherent semantic classes, such as containers or animals. For instance, a container–content rule takes a container meaning and transforms it to a contents meaning. These rules are not typically associated with any overt syntactic marker (contrast *The pot was washed* and *The pot was stirred*), so to ascertain whether a rule is being used requires the listener to judge which meaning is more plausible. The exact operations by which such rules occur is subject to debate; in particular it is not clear if shifts are lexical, syntactic, semantic or pragmatic phenomena (for discussion see Brennan & Pylkkänen, 2008; Copestake & Briscoe, 1995; Harris, Pylkkänen, McElree, & Frisson, 2008; Lapata & Lascarides, 2003; McElree, Traxler, Pickering, Seely, & Jackendoff, 2001; Miller, 1995; Murphy, 2007; Papafragou, 1996; Pustejovsky, 1995; Pylkkänen, 2008; Pylkkänen, Llinas, & Murphy, 2006; Pylkkänen & McElree, 2006, 2007). In addition, not every theory agrees that shifts require a system that is rule-based (Nunberg, 1979, 1995, 2004). But whatever the form of the theory, the child still needs to acquire a system that is productive.

While the child needs to learn which shifts to make, they also need to learn which shifts not to: Not every shift is possible. For example, although each shift in (1–3) above is licensed, shifting in the opposite direction is not possible. *DVD* has a sense similar to movie, but *movie* cannot be shifted to 'DVD the movie is on' to derive a plausible reading for (4). *The soup* is not easily shifted to its container in (5), and even though we interpret *Book* as 'Reading the book' in (3), we are unable to perform the same shift in reverse (6). Furthermore, there is limited cross-linguistic work demonstrating that the set of licensed shifts varies across languages. For instance, Kamei and Wakao (1992) argue that the producer–product shift (*Barty read Dickens*) is not licensed in Mandarin Chinese.

(4) The movie was shiny.
≠ *The DVD containing the movie was shiny.*
(5) The soup was cracked.
≠ *The pot containing the soup was cracked.*
(6) Reading the book was 200 pages.
≠ *The book that was read was 200 pages.*

1.2. Shift acquisition

How might children learn the set of licensed shifts while excluding the unlicensed ones? Because theories of lexical development (e.g., Bloom, 2000) assume words are form-meaning pairs, they cannot provide an adequate answer. If children were to learn word senses piecemeal, then each novel sense should be equally plausible, and this is clearly not the case: some novel senses (e.g., *saucer of Sasquatch*) are easily interpretable, but others (e.g., *shiny movie*) are not. Children, then, have to learn a productive system with a set of constraints on what makes a sense licensed. This means that the learner will face problems that are analogous to those encountered learning similarly generative systems, like syntax.

The foremost difficulty for any theory of syntactic development is the projection problem: avoiding the acquisition of an overly general grammar (Baker, 1979; Braine, 1971; Pinker, 1984, 1989). As an example, when learning about argument structure the child cannot simply assume that every verb undergoes passivization. *Billy owns the books* alternates with *The books are owned by Billy*, but

[1] The sasquatch, also known as Bigfoot, is an alleged, ape-like creature inhabiting the Pacific Northwest of the USA.

[2] We use the term 'shift' to describe the process underlying a range of productive lexical phenomena, which go under labels as varied as polysemy, metonymy, coercion, systematic polysemy, deferred interpretation, sense transfer and more.