Comprehending New Words Beyond Their Original Contexts

Gregory L. Murphy

When one encounters a novel word, it may be in a context that does not specify its complete meaning and possible uses in other contexts. One source of information speakers can use is how similar, familiar words are used. I describe preliminary results from an experiment in which people read novel words in one context and then had to understand them in a new use. In some cases, these uses followed patterns of familiar words. The question is whether such patterns aid in comprehension even when the word has never been encountered in that use before.

Key words: novel word, familiar word, interpretation, experiment

1. Introduction

My route to the question of how people understand novel words is through the issue of how people understand novel uses of familiar words, and, more generally, meaning extension throughout the lexicon. My starting point is the issue of *polysemy*, the phenomenon that words have multiple related meanings, called *senses*. (In homonymy, the meanings of a word are not related.) It is generally agreed that the more frequent a word is, the more polysemous it tends to be, so that the issue of polysemy arises most often for the most frequent words we use, making it an important aspect of everyday language use. A simple, concrete word like *hit* has 15 senses listed in my (abridged) dictionary as a transitive verb, 6 as an intransitive verb, and 9 as a noun. Thus, depending on whether one considers these a single word or as three different words in the lexicon, *hit* could have as many as 30 different senses, and at least 15.

Exactly how to deal with all these different senses is not completely agreed upon in linguistic theory. One approach has simply been to list all the senses much as the dictionary does. Another view takes the opposite approach, to list a single meaning for the word and to let the other senses be derived in context (Ruhl 1989). For example, one might not list that *hit* has the sense of 'to affect strongly', but when encountering a sentence like *The town was hit by the floods*, this sense can be derived from the basic meaning of the verb when interpreted in context. Finally, a number of writers have taken a middle approach, arguing that reducing the meaning of a polysemous word to a single meaning is not in general possible, but that listing every conceivable sense is excessive (e.g., Rice 1992; Tyler & Evans 2001).

This debate (reviewed in Murphy, in press) raises the issue of how people might derive the senses of a polysemous word, if all the senses are not stored in their lexicons. For example, imagine that you have stored the meaning of the word *gnu* as, roughly, an African mammal. Because you haven't seen it used in different ways (yet), you have not stored other possible senses. But some day, you might encounter this word used in a different way. One new sense might be the entire class of gnus, as in *The gnu is found throughout southern Africa*. Another sense might refer to the meat of the gnu, as in *Gnu tastes much like gazelle*. To understand such utterances, you will first have to try to use the sense you have stored. However, that sense will not fit these contexts very well (an individual gnu would not be found throughout Africa; one does not taste a whole animal). Therefore, listeners may have to construct a novel sense from this word. For example, one normally talks about the taste of food, rather than arbitrary objects. Therefore, *Gnu tastes like...* probably refers to gnu when it

is prepared as food, and since such food in our culture is predominantly the muscle flesh of an animal, here *Gnu* probably means the cooked meat from a gnu (and not its hair, for example).

The inference process I just described seems rather tenuous and perhaps unlikely to happen in real time. However, one factor that makes such inferences easier is that they often follow certain patterns. For example, words referring to animals often refer to the edible meat of that animal. Very common examples include *salmon, chicken, turkey, tuna, skate, buffalo, duck, goose,* and *eel*. This pattern is productive, in that one can refer to the meat of a new animal just by using the animal's name: *Have you ever tried pacific bonefish?*

The existence of this pattern may make it easier to understand such novel uses of words. Indeed, with novel words that refer to new forms of technology being introduced to the language on a daily basis, use of such patterns of polysemy is probably essential. For example, when MP3 files containing music were invented, people learned about *MP3* as a kind of format for encoding music. However, the term then became extended to refer to a single example of a musical selection, as in *I illegally downloaded some MP3s*. Here the word does not refer to the format but rather to something that is represented in that format. Finally, the word is (in certain circles) extended to describe the content of the material represented in that format, as in *That web site has some really good MP3s*. A *good MP3* in this context refers to a good song, not a selection that is particularly well encoded by the MP3 format.

This set of uses is paralleled by many similar words that refer to objects that contain representations, like *book, film, magazine, photograph,* and *CD*. The word *book* can refer to a certain format (e.g., in distinction to *scroll*), to an object (*a damaged book*), and to the content of the object (*a superficial book*). The existence of this pattern of meaning extensions may be a factor in the extensions of a new word like *MP3*.

In the linguistic literature, the possibility of using such patterns (often considered *lexical rules*) is generally contrasted with the listing of specific senses in the mental lexicon. For example, if you have the sense 'the content of a printed volume' as an entry in your lexicon for *book*, then you presumably do not need to use the pattern to derive that sense. And conversely, if you can use the pattern to derive the sense, then you don't need to list it. A number of authors have proposed that only idiosyncratic senses need to be represented in the lexicon, since the others can be derived by rule (e.g., Caramazza & Grober 1976; Fellbaum 2000). However, this argument relies on an assumption of parsimony that may not be psychologically valid. That is, even though a sense can be derived from a rule, it is possible that through frequent exposure and production that it becomes listed in the mental lexicon. It is hard for me to imagine, for example, that people only represent the sense of *book* indicating a physical object and do not represent the extremely frequent sense indicating the content (e.g., *interesting book, terrible book, badly-written book, engrossing book*, etc.), even though this sense could be derived by virtue of the pattern. After a few thousand exposures to the word *book* used this way, it is difficult to believe that this sense is not learned and listed in the mental lexicon. Elsewhere, I have provided evidence that argues against this derivation account for at least some senses (Klein & Murphy 2001). However, it is very difficult to get evidence for how different senses of a word are processed, because it is impossible to know for any particular speaker whether a given sense is or is not explicitly represented in the lexicon, or even whether the sense is novel. Without knowing the speaker's history, we can't tell whether he or she has encountered a phrase like *overcooked gnu*. If not, then the word is

presumably not listed in the lexicon in the sense of 'meat of the gnu'. If the person has encountered such phrases, then possibly that sense has been stored in lexical memory.

This is the point at which the study of novel uses of words intersects with the study of novel words. In an earlier study (Murphy 1997), I realized that one way to test the use of polysemic patterns was to use novel words that readers had not encountered before. This technique involves introducing a new kind of object with a new name. We can be sure that people have not learned these words prior to the experiment, since they are made up. Furthermore, as explained below, we can control what senses people encounter the words in. This allows us to test people's comprehension of words when they are used in novel ways, and in particular to compare the comprehension of word senses that do and do not follow familiar patterns. Part of learning a new word, like *MP3*, is to learn all its different senses. To what degree is such learning based on experience, and to what degree is it based on rules of meaning?

The earlier study presented new words to people in novel senses (see below). I asked people to rate how acceptable the use of the word in this novel sense was, and the results were intriguing (see Murphy 1997, for details). However, I believe that this kind of metalinguistic judgment about word uses is limited. The real question is whether lexical rules or patterns are actually used during normal language comprehension, and asking people to explicitly evaluate the acceptability of a word use is not a direct measure of that comprehension process. I here report preliminary results from new experiments that investigate language comprehension more directly.

2. The Basic Experimental Paradigm

We first (I had help from NYU undergraduates Aravinda Seshadri and Rebecca Singer) developed a 5800-word long science fiction story. The story took place on an alien planet, which allowed us to introduce a number of novel objects and events, such as an entertainment presentation (like a movie), a vehicle, new animals, a new plant, and various pieces of novel technology. All such entities had a name on this planet, and each one was introduced in a basic or "core" sense. For example, the name for the plant or animal referred to a particular example of the whole organism. There were 18 such test words. No particular emphasis was laid on these new words, which were introduced as part of the action of the story.

For example, one of the words was *jalu*, referring to a vibrantly blue plant that grows everywhere on the planet, and so was proverbial for being plentiful. Reference to the core meaning of the word was then followed by two possible extended senses. One of them was a substance extracted from the plant, which was used to treat hair. This is a familiar pattern of polysemy, in which the name for a plant is extended to include a significant substance derived from the plant (e.g., *tea, coffee, oak, maple, cotton, corn*). Half of the subjects read the word *jalu* used to refer to such a substance: "...the process of making the plant into bottles of jalu was so inexpensive...," and half read a nearly identical sentence that did not use the word in this way: "the process of making the plant into bottles of product was so inexpensive..." Thus, after reading the story, all the readers had read about the plant itself and heard it called *jalu*. All readers also read that the plant's substance was bottled, but only half heard this stuff actually called *jalu*.

The other half of the readers heard the word used to apply to an *unconventional* polysemic extension. As mentioned above, the jalu plant was proverbial for its ubiquity. Therefore, the word was extended as an adjective referring to plentifulness: "...interesting ideas, of which she had an equally jalu amount." Although there are certainly examples of names for plants being used to refer to a particular property of the plant (e.g., *orange* becoming the name for the color of the fruit), this is not a productive and expected form of polysemy (e.g., we don't use *banana* to refer to a color or shape, or *elm* to refer to height). Half the readers encountered the word used in this unconventional sense, and half read a nearly identical sentence that omitted the target word: "interesting ideas, of which she had an equally plentiful amount."

Note that by the end of the story, all subjects have read the same content. They differ only in what words were used to refer to this content. Half of the novel words they read were extended in a conventional manner (like *jalu*-product of the jalu plant), and half were extended in an unconventional manner (like *jalu*-ubiquity). After reading the story at their own speed on paper, subjects then read a continuation on the computer. The continuation completed the story itself but more importantly also used the words in both senses. The dependent measure was the time each person took to read a sentence with the critical word. Across subjects, there were four conditions formed by the combination of two variables. First, the word was used either in a conventional or unconventional sense (in this continuation); second, that use of the word had been read earlier in the story or had not.

If people are able to use polysemic patterns to comprehend novel words, then it should be fairly easy for them to understand novel uses such as *jalu* used to refer to the substance derived from the jalu plant. Whether they actually read this use earlier in the story should not be particularly important. On the other hand, novel uses that do not follow a familiar pattern should be quite difficult to understand. But this difficulty should be mitigated by the earlier occurrence of the word in this sense. So, if people read *jalu* to refer to plentiful amounts, then this sense could be encoded in their lexicons and therefore be available when the word is encountered in a new sentence. Because the senses of the words occurred in very different sentences, we could not directly compare the reading times (RTs) of the two senses. Instead, the critical dependent measure was the interaction of the two variables: Prior exposure to the word should greatly improve the unconventional use but should have much less of an effect on the conventional use.

3. Results

In our first experiment using this design, with 30 readers, the critical interaction was in fact obtained (see Figure 1). The conventional extensions were read faster overall, but as just explained, this is not directly interpretable. What is interpretable is that prior exposure to a sense sped up reading of the unconventional senses by 251 ms, whereas prior exposure actually led to slightly slower RTs in the conventional condition.

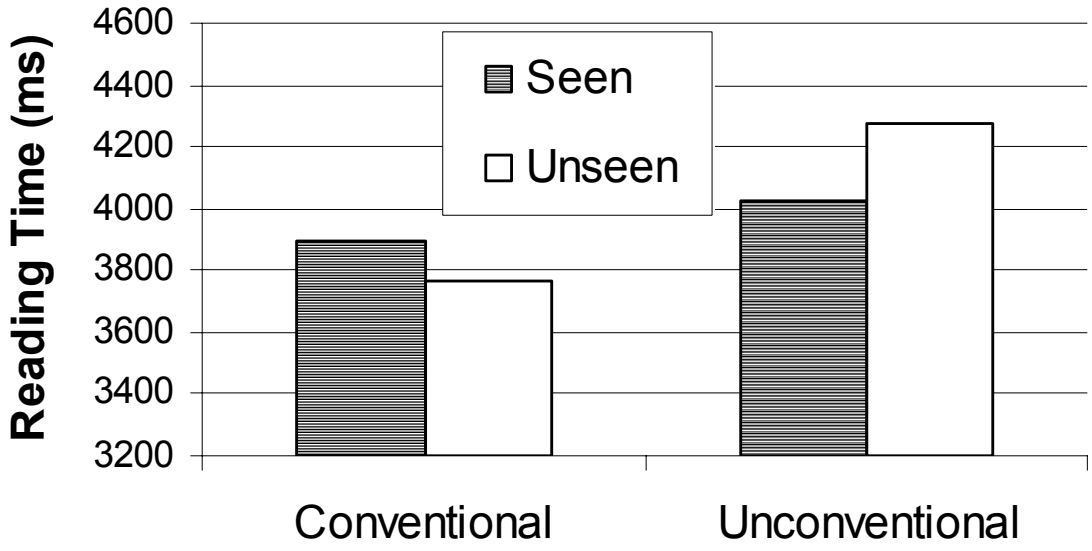

Figure 1. Reading times of sentences including novel extensions of the new words, Experiment 1. The extensions were conventional or unconventional ones, and the reader had either encountered the word used in that sense before (seen) or had not (unseen). The interaction of the two variables is reliable, $F(1, 28) = 4.19$, $p < .05$).

A second experiment is under way to improve on the first one in a number of respects. We were primarily worried that people might not remember the initial readings of the words, and in particular that 18 words were too many to encode during the brief exposure of the story. Furthermore, there was no way to be sure that readers understood the novel words, since we simply measured reading time. Therefore, Experiment 2 reduced the number of novel words from 18 to 12 and added a set of true-false questions to the story continuation, to measure whether people were carefully reading and understanding the story. Finally, at the end of the experiment, the readers were given the list of 12 novel words and were asked to write down a brief definition of them. We used this test to eliminate subjects who did not learn most of the words, probably due to lack of attention. (We were generous in giving credit—the majority of omitted words were simply left blank in the test.)

Data collection is still ongoing for this second study, but we obviously are hoping that it will replicate the pattern found in the first one.

4. Conclusions

Assuming the final results resemble the interim ones, these two experiments will provide strong evidence that people are able to extend novel words in ways that they have never heard, so long as the extension follows a familiar polysemic pattern. That is, the first time you heard the word *DVD*, it probably referred to a format and possibly a disc encoded in that format. But when you then heard someone say *I saw a great DVD last night*, you had no difficulty in understanding this (then) novel word in this novel sense. That use is just the

same as familiar uses of *book, film,* and *magazine*. Interestingly, my results do not find that reading a word in a familiar extension speeds later reading of that word in the same extension. This is somewhat surprising, as practice and frequency effects in language processing are ubiquitous (one might say *jalu*). No doubt, greater experience with a particular sense would make readers faster still to read it, even if the sense is a conventional one. But the fact that one experience did not speed reading of a new conventional extension shows just how powerful the familiar patterns are. That is, readers were apparently able to use their knowledge of such patterns to aid their understanding of a word sense the very first time they encountered it.

In summary, newly learned words can easily be extended to new senses, but such extensions are not entirely new. If they follow productive patterns already present in the language, speakers can generate and understand them easily. If they are truly new uses that do not conform to known patterns, then they must be encountered and learned, one by one. The latter finding is not very surprising, but the former provides important evidence about how new words can be easily incorporated into an existing lexicon and used in normal speech.

Acknowledgments

This research was supported by NIMH grant MH41704. Many thanks to Aravinda Seshadri and Rebecca Singer for their help in devising the stories and running the experiment. I thank the audience members and other speakers at ESSE-8 for their helpful questions and comments.

References

CARAMAZZA, Alfonso and GROBER, Ellen. 1976. Polysemy and the structure of the subjective lexicon. In C. Rameh (dd.), *Georgetown University roundtable on languages and linguistics. Semantics: Theory and application*. Washington, DC: Georgetown University Press, pp. 181-206

FELLBAUM, Christiane. 2000. Autotroponomy. In Y. Ravin & C. Leacock (Eds.), *Polysemy: Theoretical and computational approaches*. Oxford: Oxford University Press, pp. 52-67.

KLEIN, Devorah E. and MURPHY, Gregory L. 2001. The representation of polysemous words. In *Journal of Memory and Language*, Vol. 45*,* pp. 259-282.

MURPHY, Gregory L. 1997. Polysemy and the creation of new word meanings. In T. B. Ward, S. M. Smith & J. Vaid (Eds.), *Creative thought: An investigation of conceptual structures and processes*. Washington, DC: American Psychological Association, pp. 235-265.

MURPHY, Gregory L. in press. Parsimony and the psychological representation of polysemous words. In M. Rakova, G. Petho, & C. Rákosi (Eds), *The cognitive basis of polysemy*. Frankfurt am Main, Germany: Peter Lang Verlag.

RICE, Sally A. 1992. Polysemy and lexical representation: The case of three English prepositions. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum, pp. 89-94

RUHL, Charles. 1989. *On monosemy: A study in linguistic semantics.* Albany: SUNY Press.

TYLER, Andrea and EVANS, Virginia. 2001. Reconsidering prepositional polysemy networks: The case of *over*. In *Language*, Vol. 77, pp. 724-765.

Gregory L. Murphy
Professor of Psychology
Department of Psychology
New York University
6 Washington Place, Room 550
New York, NY 10003
USA
e-mail: gregory.murphy@nyu.edu