

---

# THE PSYCHOLOGY OF WORD MEANINGS

---

Edited by  
Paula J. Schwanenflugel



**THE PSYCHOLOGY  
OF WORD MEANINGS**



# THE PSYCHOLOGY OF WORD MEANINGS

Edited by

**PAULA J. SCHWANENFLUGEL**  
*University of Georgia*

A publication of the Cognitive Studies Group  
and the Institute for Behavioral Research  
at the University of Georgia



1991

LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS  
Hillsdale, New Jersey

Hove and London

Copyright © 1991 by Lawrence Erlbaum Associates, Inc.

All rights reserved. No part of this book may be reproduced in any form, by photostat, microform, retrieval system, or any other means without the prior written permission of the publisher.

Lawrence Erlbaum Associates, Inc., Publishers  
365 Broadway  
Hillsdale, New Jersey 07642

**Library of Congress Cataloging-in-Publication Data**

The Psychology of word meanings / edited by Paula J. Schwanenflugel.  
p. cm.

“Book is a product of a conference . . . jointly sponsored by the Institute of Behavioral Research, the College of Education, the Office of the Vice President for Research, and the Departments of Psychology, Educational Psychology, Language Education, and Linguistics”—Pref.

Includes bibliographical references and index.

ISBN 0-8058-0661-X

1. Meaning (Psychology)—Congresses. 2. Word recognition—  
Congresses. 3. Concepts—Congresses. I. Schwanenflugel, Paula J.  
II. University of Georgia. Institute for Behavioral Research.

BF463.M4P79 1991

401'.9—dc20

91.3444

CIP

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1



<b>3. Cross-Cultural Aspects of Word Meanings</b>	<b>71</b>
<i>Paula J. Schwanenflugel, Benjamin G. Blount, and Pei-Jung Lin</i>	
Universal or Culture-Specific Theories of Word Meaning?	72
Cross-Cultural Variations in the Structure of Word Meaning	76
Factors that Might Account for Cross-Cultural Variations in Word Meaning	80
A Cross-Cultural Conclusion	86
References	87
<b>4. The Combination of Prototype Concepts</b>	<b>91</b>
<i>James Hampton</i>	
Prototype Theory—A Formalization	92
Prototype Theory—The Evidence	94
Combining Concepts	95
Experimental Evidence	98
Noun-Noun Compounds	99
Relative Clause Conjunctions	100
Disjunction	102
Negation	103
An Inheritance Model for Composite Prototypes	105
The Composite Prototype Model	105
Evidence for the Model	108
Overextension and Dominance	109
Noncompositional Effects	112
Final Remarks—Testing Concepts to Destruction	113
Conclusion	114
References	115
<b>5. Predicating and Nonpredicating Combinations</b>	<b>117</b>
<i>Edward J. Shoben</i>	
Predicating Adjectives	118
The Modification Model	119
Complications for the Modification Model	121
Combinations and Their “Simple” Counterparts	124
Nonpredicating Adjectives	127
Relations Among Components	127
The Complexity Hypothesis	128
The Typicality Hypothesis	129
Processing Nonpredicating Combinations	132
The Relation Between Predicating and Nonpredicating Combinations	133
References	134

<b>6. Learning Word Meanings from Definitions: Problems and Potential</b>	<b>137</b>
<i>Margaret G. McKeown</i>	
A Lexicographic Context for Research on Definitions	138
Definition Problems and Solutions	142
Testing the Effectiveness of Definitions	149
Beyond Revising and Testing Definitions	152
A Final Comment	154
References	155
<b>7. Beyond the Instrumentalist Hypothesis: Some Relationships Between Word Meanings and Comprehension</b>	<b>157</b>
<i>Steven A. Stahl</i>	
Word Meaning, Orthography, and Context	161
Differences in Contexts	164
Differences Among Words	168
Differences Between Readers	169
As a Word is Learned	174
Vocabulary Knowledge and Comprehension	176
Implications for Vocabulary Instruction	179
Four Hypotheses	183
References	183
<b>8. On the Early Influence of Meaning in Word Recognition: A Review of the Literature</b>	<b>187</b>
<i>David A. Balota, F. Richard Ferraro, and Lisa T. Connor</i>	
Lexical Decision, Pronunciation, and Threshold Identification:	
The Tools of the Trade	188
The Contribution of Meaning to Word Identification: A Review of the Literature	191
General Discussion	212
Conclusion	217
References	218
<b>9. Why are Abstract Concepts Hard to Understand?</b>	<b>223</b>
<i>Paula J. Schwanenflugel</i>	
Are Abstract Words Harder to Understand than Concrete Words?	224
Theories for Why Abstract Words are Hard to Understand	238
Conclusion	246
References	247



<b>10. Interpretation of Word Meanings by the Cerebral Hemispheres: One is Not Enough</b>	<b>251</b>
<i>Christine Chiarello</i>	
Introduction	251
Cerebral Asymmetry for Language	253
Word Recognition in the Intact Hemispheres	254
Semantic Priming in Left and Right Hemispheres	255
Hemispheric Processing of Word Meanings: A Synthesis	263
Semantic Interpretation in RH-Lesioned Patients	266
Semantic Processing Without the Right Hemisphere: A Conjecture	271
Conclusion	273
References	275
<b>Author Index</b>	<b>279</b>
<b>Subject Index</b>	<b>289</b>

# Contributors

**David A. Balota** is Associate Professor of Psychology at Washington University, St. Louis, MO, USA.

**Benjamin G. Blount** is Professor of Linguistics and Anthropology and a Research Fellow of the Institute for Behavioral Research at the University of Georgia, Athens, GA, USA.

**Christine Chiarello** is Associate Professor of the Cognitive Neuroscience Program in the Psychology Department at Syracuse University, Syracuse, NY, USA.

**Lisa T. Connor** is a doctoral student in the Department of Psychology at Washington University, Saint Louis, MO, USA.

**Richard Ferraro** is Associate Research Scientist of the Department of Psychology at Washington University, Saint Louis, MO, USA.

**James A. Hampton** is Senior Lecturer of Psychology at City University, London, England, UK.

**Pei-Jung Lin** is Associate Professor of Early Childhood Education at Taipei Municipal Teachers College, Taipei, Taiwan, ROC.

**Barbara C. Malt** is Assistant Professor of Psychology at Lehigh University, Bethlehem, PA, USA.

**Margaret G. McKeown** is Research Scientist at the Learning Research and Development Center at the University of Pittsburgh, Pittsburgh, PA, USA.

**Gregory L. Murphy** is Associate Professor of Cognitive and Linguistic Sciences at Brown University, Providence, RI, USA.

**Paula J. Schwanenflugel** is Associate Professor of Educational Psychology and a Research Fellow of the Institute for Behavioral Research at the University of Georgia, Athens, GA, USA.

**Edward J. Shoben** is Professor of Psychology at the University of Illinois, Champaign, IL, USA.

**Steven A. Stahl** was a Visiting Research Scientist at the Center for the Study of Reading at the University of Illinois when this chapter was written. He is currently Associate Professor of Reading Education at the University of Georgia, Athens, GA, USA.

# Preface

This volume is the second publication of the Cognitive Studies Group and the Institute for Behavioral Research at the University of Georgia. These publications have originated from an annual conference organized by the Cognitive Studies Group of the Institute for Behavioral Research. The book is the product of a conference entitled *The Psychology of Word Meaning*, held on March 3, 1989 at the University of Georgia. The conference was jointly sponsored by the Institute for Behavioral Research, the College of Education, the Office of the Vice President for Research, and the Departments of Psychology, Educational Psychology, Language Education, and Linguistics. The participants of the conference were David Balota, Margaret McKeown, Gregory Murphy, and Edward Shoben. The other contributors to this volume were solicited through suggestions by the conference participants and organizers. I thank all the fellows of the Institute for Behavioral Research and its director, Abraham Tesser, for their support throughout this project.

**P. J. Schwanenflugel**



# An Introduction to the Psychology of Word Meaning

Paula J. Schwanenflugel  
*University of Georgia*

In the past few years, cognitive scientists have made tremendous strides in the understanding of how word meaning is represented, processed, and acquired. Simultaneously, research has progressed on a number of fronts: First, our knowledge has greatly expanded regarding fairly basic issues such as the general nature and structure of word meaning (e.g., Barsalou, 1987; Cohen & Murphy, 1984; Medin & Smith, 1984) and how such structure is reflected when word meanings are combined (Hampton, 1987; Medin & Shoben, 1988). Significant advances have also been made in our understanding of how very young children come to learn the meanings of new words (Markman, 1989) and how vocabulary learning proceeds in older children and adults (McKeown, 1985; Nagy & Anderson, 1984). Our knowledge of how word meaning is processed in language understanding has been amplified by studies of word recognition in meaningful contexts (e.g., Neely, 1991; Stanovich & West, 1983; Schwanenflugel, 1991) and studies of the processing of words with varying semantic characteristics (e.g., Balota & Chumbley, 1984; Schwanenflugel, Harnishfeger, & Stowe, 1988). Moreover, as a field, we have become much more aware of the neurological contributions to the processing of word meanings (Burgess & Simpson, 1988; Chiarello, 1988). Yet, although these various advances offer the promise of a more integrated and comprehensive understanding of the psychology of word meaning, these approaches have often seemed somewhat fragmented and isolated from one another.

The purpose of this book is to provide readers with a sense of the scope of the issues that impinge upon the psychological aspects of word meaning. I do not purport to present in this single volume a complete accounting of all the topics that might need to be included to describe the state of the art of research on word

meaning. There are, indeed, many relevant topics that have not been included that might have been. Instead, some representative, but disparate approaches have been included to elicit in the reader an appreciation for the wide range of processes that must be accounted for by a complete theory of word meaning. The goal of this book, then, is to provide an overview and organizing framework for researchers studying the psychological aspects of word meanings and related topics. I hope that the enterprise of studying word meaning has been furthered by an integrated presentation of these various topics in this book.

## OVERVIEW OF THE CHAPTERS

Specifically, this book covers five general areas in the psychology of word meaning: general theoretical issues, conceptual combinations, vocabulary development, lexical processing, and neurolinguistic issues. What you will find for each of these topics is described below.

The first three chapters are designed to address *general theoretical issues* regarding the nature and structure of word meaning. In the first chapter, Greg Murphy deals with the very difficult issue of what exactly the relation is between concepts in general and word meaning. As he points out, very often in the psychological literature we see the terms concepts and word meanings used interchangeably. He argues that this is not accidental, for many of the same processes operating in the formation and understanding of concepts in general also operate for word meaning. He concludes that word meanings are largely built out of concepts in general.

In Chapter 2, Barbara Malt proposes that examining the extensions of words provides useful information that must be accounted for by a theory of word meaning. Extension data may have distinct advantages over the intension data that is usually collected. First, a broader range of exemplars are likely to be introduced by such data than would normally be considered in studies of word meaning. Second, it avoids the need to have subjects reflect on their conscious intuitions about word meaning and, thereby, may better assess subjects' implicit semantic knowledge. The extension data that Malt presents makes it clear that, even though people may profess that verbal labels are designed to signify items with some sort of common essence, their extensions of those labels show that they do not use words to capture such common essences.

In the third chapter, Ben Blount, Peggy Lin, and I point out that language and culture have been largely ignored in studies of word meanings and concepts. We argue that culture is at the very heart of categories, concepts, and word meanings, and deserves a central focus in theories of word meanings. We review evidence that we believe shows that culture influences not only the concepts that will be acquired but also which attributes are attended to and how the world is cut up and linguistically labeled. We place the traditional Whorfian views of the relationship between language and concepts in a more modern scientific context.

The next two chapters address problems that arise when general theories of word meaning (such as prototype theory) are used to describe the process of combining the meanings of words or *conceptual combinations*. Chapter 4 by James Hampton summarizes the research on conceptual combinations, particularly the literature on noun-noun combinations. Hampton points out that a simple logical model of conceptual combinations does not work—that is, conceptual combinations do not merely involve the overlap of the extensions of each individual noun alone. School Furniture does not encompass the simple overlap of the set of schools and the set of furniture because it would exclude highly prototypical exemplars of School Furniture such as blackboards (which fit neither in the set of schools nor in the set of furniture). He describes how mechanisms might be added to the highly successful prototype theory to account for how people combine the meanings of words.

In Chapter 5, Ed Shoben describes further complications that arise for understanding conceptual combinations in his discussion of predicating and non-predicating adjectives. For predicate adjective combinations such as Large Spoon, it is not that the term Large merely increases the diagnosticity of the size attribute in extensions of the concept of Spoon. Such a modification model ignores the fact that attributes are correlated such that good examples of Large Spoons tend to be made out of wood rather than metal (as prototypical spoons tend to be). Theories of concepts and conceptual combinations must take into account these correlations. Moreover, for nonpredicating adjectives such as Servant Girl and Electric Shock, such a modification model (Osherson & Smith, 1982) doesn't work at all. He identifies 14 (!) different relations between nouns and their nonpredicating adjectives in such conceptual combinations. Clearly, a theory of conceptual combinations that takes into account the complexity of such nonpredicating adjective relations is a long way off.

Chapters 6 and 7 focus on *vocabulary learning*, in particular, learning the meanings of new words in the elementary school years. One traditional way in which children are taught to learn the meanings of new words is through the use of dictionaries. Perhaps not surprisingly, children have great difficulty learning the meanings of new words in this way. As Nagy and Herman (1987) put it, dictionaries “almost seem to be written in a secret code accessible only to those with the inside knowledge” (p. 29). In Chapter 6, Margaret McKeown argues convincingly that traditional dictionaries are abysmally poor tools for learning the meanings of words. She describes in detail many problems associated with traditional dictionaries and how they might be ameliorated. However, among their many problems is that they seem to be based on an outdated Classical View of the representation of word meaning (Smith & Medin, 1981), that is, that there is a set of necessary and sufficient features that encompasses all possible extensions of words. She argues that COBUILD dictionaries, whose definitions are based around describing prototypical instances of concept use, are vastly more effective at conveying meaning to children.

In Chapter 7, Steve Stahl notes that, for the most part, children do not learn



the meanings of new words from dictionaries. Instead, they learn them from hearing or reading words in context. However, the real focus of the chapter is on the relationship between vocabulary knowledge and reading comprehension in general. On one hand, low vocabulary knowledge will cause the child to demonstrate poor passage comprehension in certain instances. On the other hand, the child will be able to derive the meanings of new words from context on the basis of their comprehension of passages in which the new words are embedded. He discusses in detail both the conditions that influence the learning of word meanings from context as well as the conditions under which unknown words will detrimentally affect comprehension.

Chapters 8 and 9 address issues in the growing literature on the influence of word meaning on *lexical processing*. In Chapter 8, Dave Balota, Richard Ferraro, and Lisa Connor review the lexical processing literatures suggesting that word meaning may influence the processing of words prior to their full recognition. They suggest that words that are relatively concrete, polysemous, or that have highly available associates are recognized faster than words that do not have such characteristics. They invoke a *more-means-better* principle to explain the pervasive effects of word meaning on the processing of words. They suggest that adding a level referring to meaning analysis to an interactive activation model of the sort discussed by McClelland and Rumelhart (1981) would provide a mechanism for describing word meaning effects in lexical processing.

In Chapter 9, the focus is on concreteness effects this time in order to address the general question of why abstract words are harder to understand than concrete words. In that chapter, we see that concreteness effects in lexical processing, sentence processing, and vocabulary development are ubiquitous (although not unanimous). The relative ability of the dual-coding (Paivio, 1986), age-of-acquisition (Gilhooly & Gilhooly, 1979), and context availability (Schwanenflugel, Harnishfeger, & Stowe, 1988) views to account for these findings is evaluated. I conclude that a complete theory of word meaning will need to include some mechanism for describing concreteness effects.

Finally, in Chapter 10, Christine Chiarello makes the important point that, while many theories of word meanings may appear to work well as they stand, they cannot be the right model unless they also possess *neurological plausibility*. She points out that visual half-field studies consistently reveal different semantic processing profiles for the right and left hemispheres. Specifically, the right hemisphere seems to keep a wider range of potential meanings activated longer than the left. The left hemisphere is more responsible for meaning selection, inhibition, and integration. However, by keeping a greater variety of potential meanings activated in the right hemisphere, semantic reanalysis is possible should the wrong meaning be selected by the left. Together the hemispheres work to perform the rather errorfree semantic processing that people display. Currently, no cognitive model of word meaning proposed takes into account the differential contribution of each hemisphere.

## EMERGING THEMES IN THE PSYCHOLOGY OF WORD MEANING

As I was editing this volume, it became clear to me that there were a number of emerging themes that pervaded several chapters of the book. These should not be overlooked because they may represent a core of agreement and issues around which a full theory of word meaning may eventually be built. Here are some of the themes I noticed in no particular order of importance.

### Prototypes Revisited

One theme that emerges from the book is that what has become our conventional notion of prototypes (which is illustrated nicely, I think, by the description of it in Hampton's chapter) may need to be altered somewhat. In this book, it is noted that there are at least two complications for prototype theory as it currently stands: (1) the *flexibility* problem and, (2) the *relativity* problem.

One of the problems with conventional prototype theory is that prototypes don't seem flexible enough to carry off the task of semantic processing (Barsalou, 1987). As Hampton and Shoben note, once words are placed in a combinatorial context, their meanings become radically altered. Hampton proposes a solution to this problem that suggests that we need not throw the baby out with the bathwater, however. By adding mechanisms which delete low importance and contradictory attributes from the two individual prototypes of the conceptual combination, we can preserve the original prototype notion while enabling it to account for the apparent composite prototypes that emerge in conceptual combination.

The second complication for prototype theory is the problem it has with accounting for the cultural and developmental relativity of concepts. That is, the original formulation of prototype theory seemed to suggest that prototypes emerge because of the overlap of attributes among objects and events in the world. This suggests that concepts exist "out there" rather than "in the head," as Murphy puts it. However, Malt and Schwanenflugel, et al. note that this formulation cannot account for the cultural diversity in concepts that exists. Moreover, Murphy notes that children often don't acquire concepts in a domain unless they possess some sort of theory regarding how the domain works. Thus, word meaning is shaped by some sort of general cognitive determination of relevance and contrast. Murphy suggests that theories pick out concepts and attributes. Schwanenflugel et al. suggest that culture may be important.

On the other hand, we probably do not want to completely throw out the concept of prototypes. For one, its prediction of graded structures is an important one for which there is massive empirical support. Practically, as McKeown notes, concentrating on presenting prototypical instances helps us to build useful dictionaries for children. Rather, we may wish to reconsider how prototypes

come about, making them more opportunistic and flexible than they currently seem to be.

### How Much Episodic/Contextual Information Forms Part of Word Meaning?

Throughout the book, several of us have noted data that suggests that, when thinking about concepts, people bring in related, contextual information that seems irrelevant for understanding words presented in isolation. For example, Murphy notes that, while interpreting superordinate terms such as *musical instrument*, people tend to bring in information regarding related contextual information such as that they tend to be played on stage, that the violins are usually put together, etc. Schwanenflugel discusses data suggesting that, even when people are making simple lexical decisions, they tend to retrieve information regarding associated contextual information from prior knowledge. Similarly, Balota et al. note that the time taken to retrieve associates is fairly predictive of a word's lexical decision time, suggesting that perhaps such associates are also retrieved during lexical decision.

However, as Stahl notes, surely the task of developing word meaning is to decontextualize the meaning of a word from the contexts in which it first appeared. The issue here, then, is just how much information does word meaning encompass and just how decontextualized is it. Is word meaning merely the whole set of knowledge that we have regarding a concept? Or is it some subset of that knowledge?

We can find two distinct answers to this question in this volume. Murphy takes the position that there truly is a core around which word meanings are built. This core is picked out by the theories that people possess of a domain that decide which information is central and which information is not. In contrast, Balota et al. and Schwanenflugel take the position that perhaps word meanings never become truly decontextualized. In fact, Balota et al. suggest (as others have before them) that a model in which word meanings are cast as the accumulation of all a word's episodic traces is most appropriate.

### Theories of Attributes

Another theme that surfaces in several chapters is that people have theories about concepts that tell them which attributes are likely to be important to attend to and which attributes are not important. As Murphy notes, we have theories that tell us that there is probably no thing that both writes poetry and performs photosynthesis. We may have a mistaken theory that all living things breathe. As Shoben notes, these theories may enable us to distinguish which correlations among attributes are important and which ones are not. Moreover, Hampton's composite prototype model suggests that people use these theories about attribute

relations to check for the compatibility of attributes in combining concepts, making it difficult for them to know how to deal with the concepts such as “Fish that are also Birds.” Regardless of whether people’s theories are correct or incorrect, there is some evidence that they use them to form new concepts and to classify exemplars.

There are several notes of caution, however. First, there seems to be little being put forth regarding potential constraints on the theories that people construct to organize domains. Second, as Malt notes, it is often difficult to see what theories of attributes people are using to construct some categories. As she puts it, “It is difficult to imagine a theory of can-hood will include a tin tomato sauce can and exclude a tin cup, and at the same time, will include a cardboard orange juice can and exclude a cardboard ice cream carton; . . . ” (p. 65).

### The Early Availability of Word Meaning in Lexical Processing

Another theme that emerges in this book is that word meaning has its influence at a very early stage of lexical processing. The chapters by Schwanenflugel, Balota et al., and Chiarello make this very clear. Concreteness, polysemy, and associations *all* seem to have effects at the earliest point we can possibly test for them. Merely providing a letter string with a meaning (either through context or explicitly) influences how they are processed thereafter.

As Balota et al., nicely points out, these findings of early meaning influences on lexical processing suggest that models of word recognition stating that meaning is only available to the reader or listener *after* a word has been accessed (or, as they put it, after some *magic moment* in lexical access) are grossly in error. Balota et al., Chiarello, and I all agree that models of lexical processing will need to include some sort of mechanism that will enable word meaning to have an influence *prior* to the full recognition of words. Chiarello would also add that we may have to give the different hemispheres different roles in this process.

### Emphasis on Later Semantic Development

The chapters by Stahl, McKeown, and I all emphasize semantic development during the elementary school years. The usual literature on semantic development typically focuses on the kinds of intuitions that very young children have about the meanings of words. However, as is noted in Stahl’s chapter, the bulk of vocabulary acquisition does not occur prior to the elementary school years. Children come into school with about 6000 words and leave it with a vocabulary containing approximately 88,700 discrete word families (Nagy & Anderson, 1984). As my chapter on concreteness effects points out, much of this new vocabulary growth will consist of fairly abstract words.

The recognition of a large growth in vocabulary throughout the elementary

school years suggests that more emphasis needs to be placed on the kinds of processes that are used by older children to learn the meanings of new words. The vocabulary literature suggests that induction of the meaning of words from context is one important focus that research may take. Creating better dictionaries may be another.

Whatever the mechanisms behind this growth in vocabulary development, the research makes it clear that fast-mapping is not just an ability possessed by preschoolers (Carey, 1978). Even older children learn something about the meaning of a word from a single presentation in context. This fast-mapping ability appears to be one that continues to be useful throughout the elementary school years and probably beyond.

## REFERENCES

- Barsalou, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 101–140). Cambridge: Cambridge University Press.
- Balota, D. A., & Chumbley, J. I. (1984). Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage. *Journal of Experimental Psychology: Human Perception and Performance*, *10*, 340–357.
- Burgess, C., & Simpson, G. B. (1988). Cerebral hemispheric mechanisms in the retrieval of ambiguous word meanings. *Brain and Language*, *33*, 86–103.
- Carey, S. (1978). The child as a word learner. In M. Halle, J. Bresnan, & G. Miller (Eds.), *Linguistic theory and psychological reality* (pp. 264–293). Cambridge, MA: MIT Press.
- Chiarello, C. (1988). Lateralization of lexical processes in the normal brain: A review of visual half-field research. In H. A. Witaker (Ed.), *Contemporary reviews in neuropsychology* (pp. 36–76). New York: Springer-Verlag.
- Cohen, B., & Murphy, G. L. (1984). Models of concepts. *Cognitive Science*, *8*, 27–60.
- Gilhooly, K. J., & Gilhooly, M. L. (1979). Age-of-acquisition effects in lexical and episodic tasks. *Memory & Cognition*, *7*, 214–223.
- Hampton, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*, *15*, 55–71.
- Markman, E. (1989). *Categorization and naming in children*. Cambridge, MA: The MIT Press.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, *88*, 375–407.
- McKeown, M. (1985). The acquisition of word meaning from context by children of high and low ability. *Reading Research Quarterly*, *20*, 482–496.
- Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology*, *20*, 158–190.
- Medin, D. L., & Smith, E. E. (1984). Concepts and concept formation. *Annual Review of Psychology*, *35*, 113–138.
- Nagy, W. E., & Anderson, R. C. (1984). How many words are there in printed school English? *Reading Research Quarterly*, *19*, 304–330.
- Nagy, W. E., & Herman, P. A. (1987). Breadth and depth of vocabulary knowledge: Implications for acquisition and instruction. In M. G. McKeown & M. E. Curtis (Eds.), *The nature of vocabulary acquisition* (pp. 19–35). Hillsdale, NJ: Lawrence Erlbaum Associates.

- Neely, J. H. (1991). Semantic priming effects in visual word recognition. In D. Besner & G. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264–336). NJ: Lawrence Erlbaum Associates.
- Osherson, D. N., & Smith, E. E. (1982). Gradedness and conceptual conjunction. *Cognition*, *12*, 299–318.
- Paivio, A. (1986). *Mental representations: A dual coding approach*. Oxford, England: Oxford University Press.
- Schwanenflugel, P. J. (1991). Contextual constraint and lexical processing. To appear in G. B. Simpson (Ed.), *Understanding word and sentence*. Amsterdam: North Holland.
- Schwanenflugel, P. J., Harnishfeger, K. K., & Stowe, R. W. (1988). Context availability and lexical decisions for abstract and concrete words. *Journal of Memory and Language*, *27*, 488–520.
- Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- Stanovich, K. E., & West, R. F. (1983). On priming by a sentence context. *Journal of Experimental Psychology: General*, *112*, 1–36.



# 1

## Meaning and Concepts

Gregory L. Murphy  
*Brown University*

When a child learns to call only cars by the name *car*, has he or she learned the meaning of *car*, the concept of cars, or both? Could the child have learned one without the other? Imagine an adult who calls some tree *an oak* when it is really a maple. Does this person not know the correct meaning of *oak* or is he or she confused about the category of oaks? Or do these questions amount to the same thing? The goal of this chapter is to analyze the psychological representation of word meaning and to determine its relation to concepts. In carrying out this goal, we will also need to consider philosophical views on what word meaning really is.

When discussing meaning and concepts, the first problem to address is which is which. In the psychological literature, some writers seem to randomly choose one term to use, without any clear motivation for the choice. Other writers use the terms interchangeably, perhaps in the hope that at least one of them will be correct. The issue here, of course, is not to choose one term as being “correct,” but to decide what these terms mean in a theory-neutral way. By *meaning* I mean the semantic components of words, in particular. That is, meaning is the component of linguistic elements that gives them significance. This usage does not imply a commitment to any particular theory of meaning. By *concepts* I mean mental representations of coherent classes of entities. Concepts are our notions of what kinds of objects and events make up the world. These representations may or may not correspond to word meanings, and this is a large part of what we’ll be worried about in this chapter.



## SEMANTIC THEORIES

Let me start by very briefly considering one view of what meanings are—the view given to us by formal semantics (see Dowty, Wall, & Peters, 1981, for a good introduction). Formal semantics considers meaning to be a relation between words and the world. There are two parts to this relation. First, there is the *extension*, which is the set of all the objects that the word describes. So, the extension of *chair* would be all the chairs in the world—or, according to most theories, all the possible chairs in all the possible worlds. The second component of meaning is called the *intension* (not to be confused with *intention*). In everyday terms, the intension is the property that all chairs have in common—it's the chairness of chairs. One way to think of the intension is as a rule or property that can pick out the extension. Once you know what the (true) property of being a chair is, you can pick out all the chairs that you might encounter. Most theories of formal semantics use both of these components, which causes a serious problem for psychologists trying to apply these theories to psychology, because it seems very clear that the extension can't be anything that people know or manipulate. That is, people cannot know or represent all the chairs in the world, much less in all possible worlds. To this degree, then, it's very difficult to relate such formal semantic theories to psychological accounts of meaning (Cohen & Murphy, 1984).

This chapter focuses on the more psychologically relevant aspect of meaning, the intension. It seems likely that people could learn these intensions, or something like them, when they learn word meanings. And if you know the intension of a word, you may not have to know its complete extension, because the intension can be applied to actual objects to evaluate whether they belong to the extension. So, if you know what the intension of *chair* is, you can use it to pick out actual chairs. If intensions are part of word meanings, then we can make at least some connection between semantic theory and human language processing through them.

Virtually every theory of the psychological representation of word meaning proposes that people learn the intensions of words, with different theories disagreeing on just how the intensions are represented. For example, theories have claimed that people represent word meaning as semantic components (Clark, 1972; Katz & Fodor, 1963; Smith, Shoben, & Rips, 1974), as prototypes (Rosch, 1973, 1975), as nodes in a network (Collins & Quillian, 1969), as mental models (Johnson-Laird, 1983), and so on.

However, the philosopher Hilary Putnam (1973, 1975, 1988) has made compelling arguments against the assumption that people know the intensions of words, which is at the core of our psychological theories of meaning. The following example illustrates Putnam's argument. People have been buying oranges and lemons for many years, and most of them know what the difference is

between them. However, it's conceivable that a biologist could do a genetic and morphological study of lemon trees and find out that some of the varieties that we had been calling lemons are really a rather peculiar variety of orange. This may sound a little outlandish, but such re-evaluations have actually occurred. We know that whales are not fish, even though they look and act a lot like them, and people previously thought that they were fish. Biologists have made similar surprising discoveries over the years (e.g., that "glass snakes" are really a kind of lizard), so this one is not unreasonable.

In this example, it turns out that even fluent speakers of English have been misusing the language. They have been calling a large number of oranges *lemons* for years, as if they didn't know the meaning of the word *lemon*. Although word meaning may not at first glance seem to be an empirical matter, the scientist's discovery would change people's word use and understanding of what lemons really are. So, many people who speak English apparently didn't know what the word *lemon* means—although they had a mental description of lemons, this couldn't have been the true intension of *lemon*, as it didn't actually pick out only lemons (but also some oranges). The point is, of course, that this scenario could arise for almost any natural kind word—not just lemon. So, the same kind of example could be used to show that people don't really know what *horse*, *cat*, *fruit*, *water*, *canyon*, *tree*, *silver*, or *grass* mean. It's not important that scientists haven't yet disconfirmed our notion of, say, what *silver* means—the very fact that scientists *could* disconfirm it shows that there is more to the meaning than whatever is in our mental representation of the word.

The example of the lemon is somewhat realistic, in that scientific discoveries are often changing our conceptions of things that we thought we understood. Putnam (1973, 1988) has constructed a more whimsical example that makes the point even more compellingly (my description is also based on the version of Fodor, 1987). Imagine that there is a world called Twin-Earth that is identical to the Earth in every respect but one. For example, every person on Earth has a twin there who is atom for atom identical with his or her Earth counterpart. The only respect in which Twin-Earth and Earth differ is that whereas our substance gold corresponds to an element, the thing that they call "gold" on Twin-Earth is a complicated alloy, which we can call "XYZ." Superficially, however, gold and XYZ are very similar.

Now let's consider the relationship between me and the person just like me on Twin-Earth (Twin-me). Presumably, we must have the same thoughts and mental life, because we are molecule-for-molecule identical. However, Putnam argues that our meanings are not identical. For example, we might both make the statement, "I have a gold watch." However, when I say that sentence, what I mean is that I have a gold watch, but what my twin means is that he has an XYZ watch. That is, the meaning of *gold* on Earth is the element gold, but the meaning of the same word on Twin-Earth is XYZ. Since it is XYZ that they have

always called gold, XYZ is causally connected to the word *gold* in just the same way that gold on Earth is connected to the word *gold*.

Why is this a problem? Well, since Twin-me and I have identical brains and, it seems likely, identical mental representations, it seems that we should have identical meanings. However, as just pointed out, we don't. We differ in that my meaning of *gold* is not the same as my twin's. Thus, Putnam poses the following paradox: Even though we have identical mental and brain states, our meanings are not identical. For psychologists, the point is that word meaning doesn't seem to be a matter of psychology (i.e., our mental representations); the physics of Earth or Twin-Earth partly determines what the meaning is.

Putnam points out several important implications of these examples. One is that there is a *linguistic division of labor* for names describing natural kinds: We know something about the meaning of these words, but we also delegate some of the responsibility for determining their true meaning to other people—experts. To some degree, then, it's an empirical question as to exactly what a lemon is or gold is; the science of biology will tell us (eventually) what it really is to be a lemon, and we have to allocate some of the linguistic work to biologists at the same time that we're using these words.

The second implication is that people apparently don't know the intensions of words. It seems that we didn't really know the meaning of *lemon*, even though we had a mental description of lemons. And even though Twin-me and I had identical mental states, and, in particular, we had the same beliefs about gold, we had different meanings of the word *gold*. As Putnam (1973) put it in a famous quote: "Cut the pie any way you like, 'meanings' just ain't in the head!" (p. 704). Why aren't they in the head? Because there is a "right answer" to what the meaning is, and that answer is a fact about the world—not about our mental representations. So, the meaning of *lemon* is whatever lemons really are, regardless of our beliefs about them. And the meaning of *gold* depends on what gold really is on Earth (or Twin-Earth), not just on what we believe about gold.

One may feel that the Twin-Earth example is a rather *récherché* philosophical puzzle that has few implications for psychology. But its importance is in suggesting that meaning is not a matter of mental representations, and therefore that the study of meaning is not really part of psychology. And cases of meaning change, such as the lemon/orange example, show that Putnam's problem can have real consequences—namely, that people's mental representations of a word don't seem to determine what the word really means.

If we accept Putnam's conclusion, the psychology of meaning looks like a doomed endeavor. If meanings ain't in the head, there can't be any psychology of them. How is a psychologist to deal with this problem? To begin with, by ignoring it. Rather than directly answer the question of where meanings are, if they aren't in the head, I'm going to approach the question from the other side by asking just what is in the head. Rather than attacking Putnam's problem itself, I will initially confine myself to the question of psychological semantics: How do

people represent word meanings? I'll discuss this by working through a series of questions that spell out the major possibilities. Following this purely theoretical discussion is a review of empirical evidence that bears on the question of meaning representation. Finally, the chapter ends with an attempted resolution to Putnam's problem by reviewing what has been found in the head and evaluating whether it is "meaning."

## THE REPRESENTATION OF MEANING

Perhaps the first question one should ask is whether people have any knowledge at all of word meanings. Although the answer seems obviously to be "yes," there are possible arguments against it. One could argue that English is a language with a grammar and semantics regardless of whether anyone alive speaks it (Katz, 1981, expresses this view but not the psychological implication under discussion). English could be conceived of as a linguistic system apart from its speakers, just as mathematics was a valid system even before humans discovered it. To take a realistic example, there certainly is a fact of the matter as to whether a phonetic sequence is a sentence of ancient Hittite, even if there is no one alive who knows this language. The rules of the language do not have to be actually known by anyone in order to be the correct rules. Therefore, one could argue, there's no need to talk about meanings of English words being in the head, because the semantics of English exist (in some sense) independently of people.

This view of language has a certain degree of validity (though see Chomsky, 1986), but it doesn't seem to bear on the question of what speakers of a language actually know. The fact that one could imagine the semantics of English existing apart from any speakers does not entail that speakers don't know the rules of English. In contrast, people's everyday language use provides strong *prima facie* evidence that they know something about their language, or else their behavior would be inexplicable. Although Putnam's examples still must be accounted for, it is clear that there is something in the head that accounts both for performance in experiments (see Smith, 1978) and for normal language use.

Let us tentatively conclude then, that people have some kind of mental representation that controls their word use. What is the nature of these mental objects? One possibility is that meanings are a purely linguistic construction, just as relative clauses and phonemes are purely linguistic elements—someone who doesn't know any languages doesn't know about phonemes or relative clauses. On this view, meanings are strictly internal to the linguistic system; they aren't part of the general "language of thought" that we use to think about the world.

If we had such a theory, we could still explain the semantic relations among words. For example, we could take the Katz and Fodor (1963) framework of semantic markers as being such a theory. They suggested that word meanings can be represented as collections of semantic markers, such as [unmarried, adult,

male] for the meaning of *bachelor*. If we interpret the semantic markers as arbitrary linguistic symbols, then we might have a theory of meaning that was strictly internal to the linguistic system. That is, the features could be rewritten as any arbitrary strings with no effect on the theory's explanatory power. The feature [male] could just as well have been [xyzyzy] or [%3#A20!]. (By the same token, if syntacticians were to call noun phrases "LGs" instead of "NPs," this would have no effect on the explanatory power of their theories.) We could still explain semantic relations such as entailment, synonymy, and antonymy through the relations of these symbols. For example, synonymous words would have exactly the same (arbitrary) semantic markers. Opposites would be words that differ in the value of exactly one marker.

Such a theory could explain many linguistic phenomena, but it has serious problems as a psychological model. In order for people to understand how words are related to real objects in the world, they must interpret each of these components in terms of some concept that they already know. Although one could have a set of semantic markers that are purely internal to the linguistic system, these meanings couldn't refer to real objects by themselves, because they aren't connected to those objects in any way. That is, in order to identify birds with the word *bird*, the word must be connected to actual properties of birds. And in order to use the word *bird* in conversation, it must be connected to knowledge of all kinds about birds. So each semantic marker would have to be tied to a concept or subconcept that people have, so that they can connect words to the things they represent in the world. For example, the feature [xyzyzy] would have to be connected to one's concept of maleness in order to accurately label someone with the word *bachelor*. It should be obvious that this view posits a semantic system that is completely redundant with the conceptual system. That is, for every semantic marker that influences word use, there must be a conceptual element that connects the marker to the world. Why not just use conceptual markers to begin with? Or to put it another way, the semantic markers don't actually supply the meaning of the word; it's the concept it's connected to that makes [xyzyzy] mean maleness instead of redness or some other attribute.

In order to form a psychological model, then, it seems necessary to hook semantics up with concepts in some way. Perhaps the most straightforward way would be to suggest that meanings and concepts simply are the same thing (still ignoring Putnam's arguments for the moment). That is, perhaps every concept is a meaning, and every meaning is a concept. Unfortunately, this simple proposal won't work. Clearly, we have concepts that we don't have words for. (In fact, a few years ago, a book called *Sniglets* came out that invented words for concepts that we didn't yet have words for—such as *sniglet*, which is a word made up to fill such a lexical gap.) Many complex concepts have no conventional word associated with them (e.g., the concept of things to do at the beach when it's raining—see Barsalou, 1983). And it seems clear that children have many concepts that they haven't learned the words for yet (Clark, 1983). So every concept probably isn't a word meaning. Furthermore, it seems likely that some words

don't map cleanly onto a single concept, but that they may pick out a complex construction of a number of concepts (e.g., *sniglet* or *democracy*). So, every word meaning may not be exactly one concept. In short, the simple proposal that word meanings are identical to concepts is too simplistic.

If meanings aren't equal to concepts, and yet if concepts are involved, then perhaps the best way to describe their relation is to say that meanings are built out of concepts in some way. That is, a word's meaning is constructed by mapping concepts onto the semantic component of the lexicon. This is hardly an original conclusion (e.g., Clark, 1983), but by considering and eliminating the possibilities raised earlier, we can feel more confident that there is a connection between word meaning and the conceptual system. In addition, it is clear that concepts and meanings are not exactly identical and should not be treated as interchangeable.

We can't feel too confident in this conclusion, however, because a number of writers on semantics have criticized a view similar to it. Janet Fodor (1975, p. 16) and Clark and Clark (1977, ch. 11) argued against the idea that meanings are "ideas." (In particular, Fodor focused on the notion that one's ideas are mental images. I will apply her arguments more broadly.) They claimed that different people have different ideas about things, whereas they share the same word meanings. They also suggest that one's personal idea about an object changes depending on the context, one's mood, and so on, whereas word meanings should be relatively constant. In short, "ideas" are rather ephemeral, unstable things, whereas meanings are not.

One way to generalize this argument is to say that language is a kind of social convention that we generally share. We can't make up our own word meanings any more than we can make up our own syntactic rules and think that we're still speaking English. To that degree, then, meanings can't just be our individual ideas. However, to the degree that meanings must be represented by each language user, it isn't so clear that the *psychological* basis of word meaning can't be our "ideas" or concepts (although there may also be a societal basis that could require a different kind of analysis; see below). Consider the argument that people might have different concepts about things. For example, you could have a different concept of dogs than I do (Barsalou, 1987). But if this were true, you might also use the word *dog* somewhat differently than I do. If you were to say "It's a dog's life," you might intend something different than if I said it. Rather than being a problem, this is a point in favor of the psychological theory that meanings are built from concepts.<sup>1</sup> Of course, you probably wouldn't have a

---

<sup>1</sup>Actually, this issue is somewhat more complex than the discussion admits. It may be important for *linguistics* to assume that everyone in the community shares the same word meanings, even if this is not strictly true (Chomsky, 1965). If one takes psycholinguistics to be a branch of linguistics, then the first argument might well apply to it. If one takes psycholinguistics to be a theory of performance, then the argument can be dealt with as it is in the text. Unfortunately, there is still little agreement on this matter in the field.

concept of dogs that's radically different from everyone else's, because you would soon be corrected once you started using the word *dog*. Thus, although one could imagine people's concepts diverging wildly, there are strong social constraints within a language community to prevent this from happening. If there are consistent individual or dialect differences in word use, then there should be corresponding differences in concepts.

Barsalou's (1987, 1989) research on conceptual instability is of considerable relevance to the second argument, that concepts differ even within individuals. His experiments show that many different measures of category structure have little stability across subjects from the same population, across contexts, and even within individuals across test sessions with identical instructions. Barsalou argues that it is a mistake to think of one mental structure that is "the" concept. Rather, there is a wide variety of different kinds of information connected to the representation, and which set of information gets activated in any situation depends on the person's knowledge, recent experience and the current context. If this analysis is correct, then the phenomenon noted by Fodor and others may be even more widespread than introspection indicates. But is this phenomenon an argument against using concepts to represent word meanings? Barsalou's experiments all involve linguistically described categories, including those denoted by common nouns. Therefore, his results probably have implications for linguistic processes involving those words, such as lexical selection in production and disambiguation in comprehension. And one of Barsalou's tasks required subjects to provide word definitions—presumably a task central to semantic representation. This task showed the same instability as the others (Barsalou, 1989).

In short, Fodor and others who have made this argument may be right that "concepts" are changeable and context-sensitive. However, they haven't demonstrated that the mental representation of word meaning is *not* similarly changeable and context-sensitive (and there is some evidence that it is). So, this is not a decisive argument against the conceptual representation of meaning.

I should add that it seems unlikely that our long-term concepts or word meanings change very much within short periods of time. We may emphasize one part more than others at different times: I may be disgusted with dogs or happy with dogs, but even as my attitude changes, my basic concept of what dogs are doesn't change that much. For example, as I become increasingly depressed, I don't begin to think that dogs are robots or plants, or that they can fly. Their essential nature is unchanged as a function of my mood.

There's a more serious problem with a view that meanings are built out of concepts, namely that concepts may contain considerable information that does not seem to be part of word meaning (Clark & Clark, 1977, pp. 411–412). I illustrate this issue with some experiments on categorization. This research (reported in Murphy & Wisniewski, 1989a) investigated the difference between basic concepts like *lamp*, *dog*, and *car* and more general concepts (called "superordinates") like *furniture*, *animal*, and *vehicle*. In the standard categorization