

Semantic Leaps

*Frame-Shifting and
Conceptual Blending in
Meaning Construction*

SEANA COULSON

University of California, San Diego




CAMBRIDGE
UNIVERSITY PRESS

Semantic Leaps

Frame-Shifting and Conceptual Blending in Meaning Construction

Semantic Leaps explores how people combine knowledge from different domains to understand and express new ideas. Concentrating on dynamic aspects of on-line meaning construction, Coulson identifies two related sets of processes: frame-shifting and conceptual blending. Frame-shifting is semantic reanalysis in which existing elements in the contextual representation are reorganized into a new frame. Conceptual blending is a set of cognitive operations for combining partial cognitive models. By addressing linguistic phenomena often ignored in traditional meaning research, Coulson explains how processes of cross-domain mapping, frame-shifting, and conceptual blending enhance the explanatory adequacy of traditional frame-based systems for natural language processing. The focus is on how the constructive processes speakers use to assemble, link, and adapt simple cognitive models underlie a broad range of productive language behavior.

Seana Coulson is Assistant Professor in the Department of Cognitive Science at the University of California, San Diego.



Semantic Leaps

*Frame-Shifting and
Conceptual Blending in
Meaning Construction*

SEANA COULSON

University of California, San Diego

 **CAMBRIDGE**
UNIVERSITY PRESS

for Steven

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
10 Stamford Road, Oakleigh, VIC 3166, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa
<http://www.cambridge.org>

© Cambridge University Press 2001

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2001

Printed in the United States of America

Typefaces Palatino 10/13 pt. and Optima System $\text{\LaTeX} 2_{\epsilon}$ [TB]

A catalog record for this book is available from the British Library.

Library of Congress Cataloging in Publication Data

Coulson, Seana, 1966-

Semantic leaps : frame-shifting and conceptual blending in meaning construction /
Seana Coulson.

p. cm.

Based on the author's thesis (doctoral - University of California, San Diego, 1997)

ISBN 0-521-64361-9

1. Semantics - Psychological aspects. 2. Frames (Linguistics) 3. Concepts.
4. Cognitive grammar. I. Title.

P325.5.P78 C68 2000
401'.43 - dc21

00-036046

ISBN 0 521 64361 9 hardback

Contents

Acknowledgments	page xi
1 Semantic Leaps	1
1.1 Productive Language Behavior	4
1.1.1 Cats on Mats	4
1.1.2 Mats on Cats	8
1.1.3 The Centrality of Mapping	10
1.2 Background	17
1.2.1 Frame Semantics	17
1.2.2 Mental Space Theory	21
1.3 Connections	28
PART ONE. FRAME-SHIFTING	31
2 Frame-Shifting and Models of Language Processing	33
2.1 Approaches to Comprehension	36
2.1.1 Composition: A Thoughtful Wife	38
2.1.2 Semantic Indeterminacy	41
2.1.3 Sense and Nonce Sense	42
2.1.4 Frame-Shifting: Another Thoughtful Wife	44
2.1.5 Constructivist Accounts	45
2.2 Frame-Shifting in One-Line Jokes	49
2.2.1 Typical Situations	49
2.2.2 Default Values	55
2.2.3 Reinterpretation	56
2.2.4 Goal Presupposition	58
2.2.5 Nonconventional Meanings	62
2.3 Conclusions	69

3 Models of Sentential Integration	71
3.1 Frame-Shifting in Text Processing	73
3.2 Frame-Shifting in Sentential Integration	75
3.3 Frame-Based Models	83
3.3.1 Frames and Constraint Satisfaction	84
3.3.2 The Story Gestalt Model	85
3.3.3 Structured Connectionism	87
3.3.4 Space Structuring	88
4 Frame-Shifting and the Brain	92
4.1 Neuropsychology	94
4.2 Event-Related Brain Potentials	98
4.2.1 Lexical Versus "Post"-Lexical Processing	100
4.2.2 What's in a Word?	104
4.2.3 ERP Indices of Frame-Shifting	106
4.2.4 Summary	113
PART TWO. CONCEPTUAL BLENDING	115
II.1 Trashcan Basketball	115
II.2 Blending and Integration Networks	118
II.2.1 Single Framing Networks	120
II.2.2 Frame Networks	120
II.2.3 Blending Processes	122
II.2.4 Outline	123
5 Conceptual Blending in Modified Noun Phrases	125
5.1 Nominal Compounds	126
5.1.1 Conceptual Blending and Nominal Compounds	128
5.2 Predicating and Nonpredicating Adjectives	133
5.2.1 Predication and Noun Modification	138
5.2.2 Conceptual Blending and Predicating Modifiers	141
5.3 Privative Adjectives	144
5.3.1 The Sense Generation Model	144
5.3.2 Fake Guns	146
5.3.3 Stone Lions	150
5.3.4 Land Yachts and Luxury Cars	152
5.3.5 Summary	156
5.4 Conclusions	157
5.4.1 Analyzability and Compositionality	159

6 Conceptual Blending in Metaphor and Analogy	162
6.1 Blending in Metaphors	165
6.1.1 Cross-Domain Mapping in One-Sided Networks	167
6.1.2 Two-Sided Networks	168
6.1.3 Digging Deeper	170
6.1.4 Born on Third Base	172
6.2 Retrospective Projections	178
6.2.1 The Menendez Brothers Virus	179
6.2.2 The Menendez Brothers Virus II	183
6.2.3 The Right to Life Virus	185
6.2.4 Heidi Fleiss and Her Counterpart	191
6.2.5 Summary	196
6.3 Conclusion	197
6.3.1 Functions of a Blend	200
7 Counterfactual Conditionals	203
7.1 Truth, Acceptability, and Conditionals	204
7.2 Analogy and Identity	207
7.3 Frame-Shifting and Scalar Reasoning	212
7.3.1 Frame-Shifting	212
7.3.2 Scalar Reasoning	215
7.4 Counterfactual Selves	216
7.4.1 Knowledge Blends	217
7.5 Conclusions	219
PART THREE. APPLICATIONS: BLENDING, FRAMING, AND BLAMING	223
8 Framing in Moral Discourse	227
8.1 Cultural Models of Action, Responsibility, and Punishment	228
8.2 Framing and the Morality of Abortion	229
8.2.1 Choice, Trauma, and the Rape Exception	230
8.2.2 Blaming the Victim	233
8.2.3 From Victims to Criminals and Back	236
8.2.4 Testimony Against the Rape Exception	238
8.2.5 Why Punish Me?	241
8.3 Mere Semantics	243

	Contents
9 Frame-Shifting and Scalar Implicature	246
9.1 Scalar Implicature	247
9.1.1 Pragmatic Scales and Rhetorical Strategies	251
9.2 Return to the Rape Exception	253
9.2.1 Linking Scales	254
9.2.2 Breaking the Links	255
9.2.3 Dialogue	258
9.3 Conclusion	264
10 The Space Structuring Model	267
10.1 My Dinner with Rodney	268
10.2 Turning the Cat on Its Head	278
10.2.1 Connections	281
Bibliography	283
Index	297

Acknowledgments

This book grew out of my 1997 Ph.D. thesis at the University of California, San Diego (UCSD). An interdisciplinary blend of topics and perspectives, in many ways it is a product of UCSD's Cognitive Science Department, where the hallways are filled with talk about everything from amnesia to aplysia, bumblebees to back-prop, the elegance of X-bar theory to the tractability of XOR. But perhaps more than anything else, it is a product of years of discussion with Gilles Fauconnier. Of course, most of these discussions were rather one-sided as Gilles barreled through the history of pragmatics, opined about the current state of cognitive science, or dove head-first into a new example. During these conversations, I did my best to scribble notes, in the hopes that I might later be able to re-create the feverish train of thought, and perhaps even decipher what Gilles meant. Over the course of time, I came to think about meaning in a whole new way.

Like many cognitive scientists, I had assumed that meaning was something speakers *compile* from linguistic input. The interesting questions in my mind were how contextual knowledge and pragmatic functions affected mechanical parsing operations and the activation of word meanings. However, in looking at actual examples of language use, it soon became clear that the impact of any given utterance could go well beyond its linguistically assembled meaning. In particular, in my 1992 study examining informants' discourse about the morality of abortion, I found that people's utterances draw on a vast set of background information, cultural knowledge, societal attitudes, covert analogies, and mappings between past, present, and future selves. While grammatical knowledge is often necessary to understand utterances in these debates, it is far from sufficient. As a result, I began to see that the apparently

trivial question of how speakers assemble *utterance* meaning was not simple – and, indeed, poses a profound and interesting problem.

My own personal Copernican revolution came with the understanding that utterance meaning is not *in* the speech signal, but actively constructed by speakers in response to linguistic and nonlinguistic cues. Like the paleontologist, the speaker has the task of combining different sorts of information to derive the overall meaning of the discourse event, to exploit her imaginative capacities and derive the life of the organism from its grammatical bones. I grew to realize that, rather than compiling meanings, people were using linguistic information to help them assemble cognitive models of the discourse event. Consequently, I set out to understand how people use contextual and background knowledge in the on-line construction of meaning. This book outlines what I have learned so far.

Traditionally, when approaching this issue, cognitive scientists have appealed to the idea of a frame, a hierarchically organized data structure for representing stereotypical knowledge. Unfortunately, the static nature of traditional frame-based systems stands in marked contrast to speakers' flexible use of language. To reconcile this tension, I point to two processes, *frame-shifting* and *conceptual blending*, intended to help solve problems faced by frame-based systems for reasoning and natural language processing. Frame-shifting is semantic reorganization that occurs when incoming information is inconsistent with an initial interpretation, and conceptual blending is a set of cognitive operations for combining frames from different domains.

I can lay claim to neither process. Though the term *frame-shifting* is mine, the phenomenon has been noted by many before, including Arthur Koestler, Marvin Minsky, and David Rumelhart. Conceptual blending theory, also known as *conceptual integration* theory, owes its origin to Mark Turner and Gilles Fauconnier. I have had the great fortune of being in contact with this dynamic duo as they developed the theory, and my own work has benefited greatly from our conversations and email correspondence. Similarly, my ideas about blending have profited from many discussions with Nili Mandelblit about the ubiquity of conceptual blending operations in the context of machine translation. I would like to thank them all.

One of my goals in *Semantic Leaps* is to supplement classic concepts from artificial intelligence with new ideas from cognitive semantics. Moreover, I want to use these ideas to develop models germane to psychologists and cognitive neuroscientists interested in language

processing. In my attempts to translate the abstract concerns of cognitive semantics into psychologically and neurologically plausible models, I received a great deal of assistance from my colleagues in and around the Kutas Lab at UCSD and the Van Petten ERP Laboratory at the University of Arizona. I would like to thank Marta Kutas, Robert Kluender, Jonathan King, Kara Federmeier, and Cyma Van Petten for sharing their knowledge of cognitive neuroscience. I have found our many meetings, arguments, and hours of speculative theorizing to be immensely stimulating. Special thanks are due to Marta and Cyma for their attempts to make an experimentalist out of me.

Thanks also to John Batali, Per Aage Brandt, Aaron Cicourel, Roy D'Andrade, Jonathan Folstein, Deborah Forster, Joe Grady, George Lakoff, Paul Maglio, Teenie Matlock, Todd Oakley, Sven Ostergaard, Tim Rohrer, and Eve Sweetser for their encouragement and advice. My original research was supported by a predoctoral fellowship from the McDonnell-Pew Center for Cognitive Neuroscience in San Diego. Subsequent work was supported by postdoctoral awards from the McDonnell-Pew Center for Cognitive Neuroscience in Tucson and an NIH National Research Service Award 1 F32 DC00355-01.

Lastly, I would like to thank my family for their love, support, and good humor.

CHAPTER ONE

Semantic Leaps

"One small step for [a] man, one giant leap for mankind."

– Neil Armstrong

Uttered by Neil Armstrong just before he set foot on the moon, this quote demonstrates how the same action can be interpreted in radically different ways. In this case, Armstrong alludes to a literal construal of his action and a more figurative one. Although it was a short distance from the space probe to the moon's surface, Armstrong's step would not have been possible without the combined efforts of those on the Apollo project to conceive and implement the mission. As such, Armstrong's journey to the moon was a sign of the scientific progress made in the twentieth century. The phrase "one giant leap" is not meant to be understood as a physical action, but rather as a metaphor for the scientific progress that had made the Apollo project possible.

Armstrong's action can also be interpreted metonymically. *Metonymy* is a figure of speech in which an object is referred to by one of its attributes, or by something with which it is commonly associated. In this case, Armstrong employs part-whole metonymy in which the whole species ("mankind") is evoked by one of its members. Because the journey represents the results of centuries of general scientific progress, and more specifically the efforts of the thousands of engineers, scientists, and bureaucrats involved in the Apollo program, Armstrong's voyage has been conceptualized as a public voyage on behalf of all earth-bound humanity. In the metonymic mapping, the leap corresponds to the entire journey from earth to the moon, which culminates in the final step.

The different possible interpretations of Armstrong's step provide a good illustration of how the significance of an action is dramatically

altered by a change in background assumptions. Like the interpretation of an action such as stepping onto the moon, interpreting the meaning and significance of natural language utterances depends crucially on contextual factors and background knowledge. Armstrong's utterance, for example, would be interpreted differently if he said it just before stepping into the shower. We shall see that for action and language alike, interpretation depends on the nature and scope of background assumptions.

In an effort to concentrate on tractable problems, researchers in semantics have traditionally focused on how to specify the literal meaning of phrases such as "one small step," thereby avoiding the more difficult problem of figurative meaning (see, e.g., Barwise & Etchemendy, 1989). However, more recent developments suggest the emphasis on these sorts of expressions may have been somewhat misplaced. For example, previous investigators have sought the locus of novelty in language and thought in combinatorial processes. And, while combinatorial processes are indispensable for explaining speakers' productive capacities, they do not provide a full account of novel language use. Moreover, the focus on expressions whose truth values can easily be specified has caused many investigators to overlook the constitutive role of context in meaning construction.

In this book, I locate speaker productivity in the comprehension mechanisms underlying semantic leaps – natural language constructions that yield nonobvious meanings. Characterizing the human ability to make semantic leaps will turn out to have implications for how speakers represent, organize, and use their knowledge in the production and interpretation of language. Just as Armstrong's small step relied on the heroic efforts of the members of the Apollo project, we will see how the interpretation of natural language recruits an elaborate set of meaning construction processes.

Semantic leaps is not a technical term, but, rather, a family of interesting natural language phenomena. It includes all sorts of nonstandard meanings absent from dictionaries and, typically, not computable by traditional parsers. Leaps include things such as metaphoric and metonymic expressions, hyperbole, understatement, and sarcastic quips. They also include things such as innuendo, subtle accusations, and the private meanings that can arise when people live or work closely together. Many leaps are necessary because of the way we deploy background knowledge in meaning construction. For example, until given information to the contrary, speakers tend to assume the pool in (1) is filled with water.

(1) Everyone had so much fun diving from the tree into the swimming pool, we decided to put in a little water.

Upon learning there *was* no water in the pool we do a double take, imagining the scenario in which everyone has fun diving into an empty pool. This reanalysis process is called *frame-shifting* and is the topic of Part I.

Part II deals with the sorts of leaps needed to combine information and create new concepts. In particular, Chapters 5 through 7 argue that information integration recruits *conceptual blending*, a set of noncompositional processes in which the imaginative capacities of meaning construction are invoked to produce emergent structure (Fauconnier & Turner, 1994; 1998). For instance, Chapter 5 addresses conceptual blending coded by modified noun phrases. While some examples, like "brown cow," and "red apple," don't involve much of a leap, others, like "caffeine headache," and "topless bar," rely extensively on background knowledge about the domains in question. Moreover, understanding the difference between "fake guns" and "real guns," or "stone lions" and "real lions," involves recruitment of knowledge about human practices such as deception and representation.

Some leaps are made possible by the human ability to conceptualize one thing in terms of another. For example, in (2), success in life has been conceptualized analogically in terms of success in baseball.

(2) He's a guy who was born on third base and thinks he hit a triple.

However, as if the leap from baseball to success weren't enough, note that the game in (2) has a woman giving birth in the ballpark. Chapter 6 considers this and other instances of metaphor and analogy that seem to involve a blend of elements and structure from the two domains that are being conceptualized.

Moreover, a similar blending process will be seen to operate in examples like (3).

(3) I wouldn't go out with you if you were the last man on earth.

Surely, "I won't go out with you," would suffice to make the speaker's point. However, by positing this armageddonlike scenario, she manages to drive that point home. Chapter 7 considers this and other counterfactual leaps in which speakers combine knowledge about the current state of affairs with knowledge necessary to understand an alternative reality. Overall, examples point to an important source of productive

language behavior that, until recently, has largely escaped the notice of researchers interested in meaning.

Part III addresses instances of frame-shifting and conceptual blending in real-life examples of moral discourse. Chapter 8 explores how speakers use terms to evoke culturally shared knowledge structures that can influence the way we evaluate the world around us. For example, the speaker in (4) draws on cultural understandings of action and responsibility in order to frame the rape victim's pregnancy as punishment.

(4) This bill is drawn so narrowly that it would punitively and without compassion further harm an Idaho woman who may find herself in the horrible, unthinkable position of confronting a pregnancy that resulted from rape or incest.

In contrast, the speaker in (5) presents an alternative framing.

(5) The unwanted pregnancy flows biologically from the sexual act, but not morally from it. ... Even degradation, shame, and emotional disruption are not the moral equivalent of life.

Moreover, (5) represents the use of another sort of culturally based knowledge structure, the *pragmatic scale*. Discussed in more detail in Chapter 9, pragmatic scales consist of objects or events ordered along a relevant semantic dimension. For example, the speaker in (5) evokes a scale of potentially acceptable reasons for abortion. We shall see how rejecting particular pragmatic scales can result in frame-shifting.

1.1 PRODUCTIVE LANGUAGE BEHAVIOR

1.1.1 Cats on Mats

Philosophers and linguists have each been impressed by different things about human language competence, and semantics has been shaped by both sorts of interests. For philosophers, the interesting thing about language is its *intentionality* or aboutness. How is it that an arbitrary set of symbols can represent things in the world? Linguists, on the other hand, are struck by the fact that there are virtually no limits to what competent speakers can say. We can understand sentences we've never heard before and can generate a potentially infinite number of novel ones.

These two issues, though potentially distinct, turn out to be related. The philosopher's problem of how "cat" can represent a cat has led to an

emphasis on truth and reference. The linguist's problem of how a finite brain can process a potentially infinite number of meanings has led to an emphasis on compositionality. A language is compositional if the meaning of a complex expression is systematically related to the meanings of its constituents (Frege, 1970; 1892). The interests of philosophers and linguists are complementary because semantics provides algorithms that show how the truth of a complex expression depends on the truth of its components. A good example is propositional logic, in which the truth or falsity of any well-formed formula can be systematically determined from the truth values of its component propositions and the truth tables for the connectives.

In keeping with the dual tasks of understanding compositionality and intentionality, part of what a semantic theory is supposed to account for is the relationship between the meaning of a sentence and facts about the world that support the truth of the proposition or propositions expressed by that sentence. Semanticists assume that to understand the meaning of a sentence in the indicative mode is to understand the kind of a world in which it could be used to make a true claim. Thus the main goal of semantics is to provide truth conditions for sentences.

On an objectivist account, understanding the meaning of an utterance such as (6) involves the ability to specify the conditions that make the propositions stated by the utterance true.

(6) The cat is on the mat.

Moreover, the truth conditions for a complex expression are determined systematically from those of simpler expressions. The lexical semanticist accounts for the meanings of individual words, and the compositional semanticist provides an account of how the meanings of words are composed into the meanings of their constituents, and how constituents are composed to yield the meaning of the sentence as a whole.

An important component of this account is that sentences have an independent level of representation, syntactic structure, which is used to compute an unambiguous semantic representation. Intuitively, the meaning of (6) would be derived by combining syntactic and semantic information about the components of the sentence. A "cat" is a carnivorous mammal long domesticated by humans and kept as a pet. A "mat" is a piece of coarse, woven, or plaited fabric used as a floor covering or a support. The locative relation "on" indicates that the external argument (in this case, "the cat") is in a position in contact with and supported by the top surface of the object (in this case, "the mat").

Accordingly, understanding the meaning of a sentence such as (6) involves the ability to specify the conditions that make the propositions stated by the utterance true. A semantic account should be able to differentiate the meanings of (6) and (7). Clearly, the state of the world (or worlds) in which (6) is true differs from that in which (7) is true.

(7) The mat is on the cat.

Further, this state differs in predictable ways concerning the spatial relationships between particular cats and particular mats. Formal semantics captures the intuition that anyone who claims to understand the meanings of (6) and (7) ought to be able to discriminate between a world in which (6) is true and a world in which (7) is true.

Of course, a real semanticist would not offer such a crude description as this as a semantic account. To relate words to the world, the objectivist semanticist (actually an *extensional semanticist*) employs set theory to model properties with sets of objects that possess those properties, and relations as ordered pairs (or triples, or quadruples, etc.) of related objects. Semantics provides a set of functions that convert constituents into sets that can be tested for truth in some model.

A semantic account can convert the English representation into an unambiguous semantic representation such as might be provided by predicate calculus. For example, for (8) we might derive the following representation.

(8) A cat is on a mat.

$$(\exists x)(\exists y)((Cat(x) \& Mat(y)) \& On(x,y))$$

The semantic representation is then tested for truth in a set theoretical model such as the following:

$$\{cats\} = \{Cat1, Cat2, Cat3\}$$

$$\{hats\} = \{Hat1, Hat2\}$$

$$\{mats\} = \{Mat1, Mat2, Mat3\}$$

$$\{flies\} = \{Fly1, Fly2\}$$

$$\{on\} = \{<Cat1, Mat3>, <Hat1, Cat1>, <Fly1, Hat1>\}$$

In this model, the world consists of cats, hats, mats, and flies. To find the truth value for (8), we first must find the truth value for (9) and (10).

(9) $(\exists x)Cat(x)$ There is a cat.

(10) $(\exists y)Mat(y)$ There is a mat.

To do so, we first check our model to see whether the set of cats, {cats}, contains any members, and second whether the set of mats, {mats}, contains any members. Because both of these return true, we can go on to see whether any ordered pairs consisting of a cat and a mat are in the set {on}. Thus in this model, (8) is true but (11) is not.

(11) A mat is on a cat.

The task of the semanticist is to develop ways of transforming vague and ambiguous natural language statements into one or more unambiguous semantic representations, and to develop procedures for testing those representations for truth in a model. Adequacy of an account depends on the extent to which the output of the formalism satisfies our intuitions about the truth conditions and entailments of the original natural language sentence. Note that here I've developed only a cartoon-like explanation of how the semanticist operates. The reader interested in a thorough introduction to modern formal semantics should consult Dowty, Wall, & Peters (1981).

In any case, the formal approach to meaning has directly affected influential cognitive scientists such as Phillip Johnson-Laird, whose theory of mental models (Johnson-Laird, 1983) attempts to integrate model theoretic semantics with a theory of on-line language comprehension. Clearly, the demands of psychological plausibility preclude any straightforward mapping between formal semantics and language understanding: It is unlikely, for example, that people represent infinite sets. Rather, Johnson-Laird suggests that people construct models with individuals and *finite* sets of individuals. Johnson-Laird's mental models resemble the models of semantics in two ways: First, they represent the world in a structural way; second, one can ask whether a given sentence can be true in one of Johnson-Laird's models (Barwise & Etchemendy, 1989).

And, although he finds his belief problematic, Johnson-Laird (1989: 578) remains committed to the tenet of compositionality:

A major problem confronting the present theory is to reconcile two important constraints on the process of comprehension. On the one hand, information from an utterance is integrated into the existing model as a function of the referential links, if any, between the utterance and the model; on the other hand, the interpretation of the sense of a sentence almost certainly depends on combining the senses of its constituent words according to the syntactic relations between them.

In fact, most semanticists working today have abandoned the idea that natural language is fully compositional. However, the assumptions

of the paradigm are deeply ingrained in cognitive science. Most notably, objectivist semantics has indirectly affected a huge number of researchers who have assumed that the productive nature of language use reflects the operation of something isomorphic to a compositional system (see Fodor & Pylyshyn, 1988, for a manifestolike review of the arguments for this position). The following sections show that this assumption – that compositionality in some form or other *must* obtain – has resulted in an overemphasis on combinatorial aspects of language comprehension and a corresponding underemphasis on the creative *noncompositional* mechanisms that also contribute to on-line meaning construction.

1.1.2 Mats on Cats

In common parlance, the term “meaning” applies to both linguistic expressions and particular utterances of those expressions. In contrast, semanticists reserve “meaning” for the expressions themselves. An expression such as (6) has a *meaning* that is independent of any particular context in which it might be uttered, while an utterance of this expression has a *content* that depends on things such as the particular cat to which the speaker is referring. One might hope for a semantic theory to relate expression meaning to utterance content. However, I shall argue that we should tackle the opposite problem: how meaning construction processes that generate utterance meanings might also explain why we think of words, phrases, and sentences as having meaning that is independent of any particular context.

Perhaps the main reason for shifting our focus from expressions to utterances is that meaning is never fully context-independent. Searle (1979), for instance, provides a number of examples that argue against the idea that the literal meaning of a sentence is context-independent. Rather, the literal meaning can be applied only relative to particular contextual assumptions. As we have seen, on the traditional account, answering the question of whether the cat is on the mat requires establishing which cat and which mat are being referred to, and determining whether the correct spatial relationship obtains. However, Searle (1979: 122) points to the case in which the cat in (6) is in outer space.

... suppose that the cat and the mat are in exactly the relations depicted only they are both floating freely in outer space, perhaps outside the Milky Way galaxy altogether. In such a situation the scene would be just as well depicted if

1.1 Productive Language Behavior

we turned the paper on edge or upside down since there is no gravitational field relative to which one is above the other. Is the cat still on the mat? And was the earth's gravitational field one of the things depicted in our [original] drawing?

Is the cat on the mat – or is the mat on the cat? Further, consider the case in which the mat is stiff, and has been stuck into the floor at an acute angle. Searle urges us to imagine that the cat is drugged, and lying on the top edge of the mat. Is the cat on the mat? Alternatively, the cat might be in the normal configuration but suspended by wires so that she places no pressure on the mat. Again, is the cat on the mat?

Although we *can* come up with an answer to the question of whether the cat is on the mat in each of these cases, the answers are far from determinate. Searle's point is precisely that we have to *come up with* an answer, rather than consulting the output of a context-invariant semantic function. While a semantic function might help to establish the truth of (6) in normal situations, these examples suggest that the function needs to be augmented in cases where normal conditions don't obtain. Thus Searle argues that giving truth conditions for an unambiguous sentence's literal meaning relies upon the assumption of the correct set of background conditions. Searle's argument thus suggests that truth conditions are not sufficient to specify all possible *sentence* meanings, let alone all *utterance* meanings.

Searle suggests that the establishment of the truth conditions for even the most transparent sentence is not context-independent, but relies on the assumption that conditions are normal. In this way, judgments of truth resemble judgments of category membership. While there is a great deal of disparity in people's truth and category-membership judgments concerning *atypical* cases, people tend to agree in their judgments of more typical cases (see Lakoff, 1987; Taylor, 1995). A similar process underlies both sorts of judgments. Context-independent meaning is an illusion based on the fact that a competent language user will create a context when none is provided.

Any given expression permits an indefinite number of interpretations that depend upon a speaker's understanding of the relevant context. Considering the immense variability in the utterance meaning of (6), Langacker (1987: 155) argues that a compositional account based on context-invariant meanings is unlikely to have much explanatory content.

Consider *The cat is on the mat*. Prototypically it describes a situation where a mat is spread out on the ground and a cat is sitting or lying on it. Already there is

indefinite variability, since the cat can be of any size, coloring, or subspecies; the mat is similarly variable; the cat can assume many different postures; and so on. But this is only the beginning. Possibly the mat is rolled up in a bundle and the cat is sitting or lying (etc.) on top of it. Maybe the operator of a slide show has just managed to project the image of a cat onto a mat being used for a makeshift screen. The sentence is appropriate in a mat factory where a worker has just finished decorating a mat with the outline of a feline. Conceivably a wrestler is holding an exhibition match with a tiger and has just succeeded in pinning its shoulders to the floor of the ring. The possibilities are obviously endless.

Pointing to the plethora of scenarios to which (6) might be applied, Langacker suggests there is a gap between the abstract specification of expression meaning and the rich content of particular utterances. By definition, there is some abstract commonality between the pet cat, the projected cat, and the tiger that licenses the use of the term "cat" in all three cases. However, Langacker seems to question whether there is any more content to this suggestion than the definition itself. Merely positing an unspecified abstract notion of cathood does little to explain it. Thus Langacker suggests the notion of context-independent meaning leaves an important component of productive language behavior unexplained.

After all, a speaker whose entire linguistic repertoire consisted of (6) could potentially use that philosophical chestnut to express an infinite number of *different* things. This property of utterances reflects human productivity just as surely as does center embedding. Moreover, while grammatical regularities often prove to be important for distinguishing between scenarios in which (6) does and does not apply, they do little to explain the immense variability in the configurations of cats on mats. Because the meanings provided by a compositional semantics leave this component of productive language behavior unexplained, they are necessarily incomplete. If we are ever to explain how speakers can relate (6) to all its variegated instantiations, we need to look beyond compositional mechanisms.

1.1.3 The Centrality of Mapping

Indeed, there are a number of problems that semanticists have assumed to be uninteresting, and subsidiary to the development of a theory of meaning, that turn out to be quite interesting for the cognitive scientist interested in a broader range of intelligent behavior. In particular, the way in which speakers construct utterance meaning on-line can potentially reveal general cognitive capacities. This section considers a

1.1 Productive Language Behavior

number of examples that demonstrate the variety of different uses to which people might put a term. I shall argue that these uses go beyond a straightforward relationship between words and the world. They are too complicated to be captured by objective criteria and will necessarily involve an appeal to this facet of the language user's cognitive capacities. In particular, word use reflects knowledge of typical situations, an ability to recognize abstract commonalities between domains, and an ability to conceptualize one domain in terms of another.

Consider the differences in the interpretation of "ball" in (12)–(14).

(12) Shaq dribbled past his defender and dunked the ball easily.

(13) Joe kicked the ball and smiled as it sailed through the uprights.

(14) Sammy hit the ball right out of the park.

In (12), the "ball" is a small leather sphere, orange or brown, filled with air. In (13), the "ball" is a brown, three-dimensional ellipsoid. In (14), the "ball" is a small white sphere, filled with yarn and covered with stitched cowhide. Clearly (12)–(14) each suggest a different game, and the speaker's understanding of "ball" varies as a function of the game in which it's embedded.

This is not to say that the uses of "ball" in (12)–(14) cannot be subsumed under a single abstract meaning. However, the existence of an abstract meaning does not help very much in explaining why the competent speaker interprets "ball" differently in (12), (13), and (14). It's not just that the reader *can* infer the presence of a basketball in (12), a football in (13), or a baseball in (14). She will automatically do so, and would presumably be surprised if (12) turned out to be a football, (13) a baseball, and (14) a basketball. Roth & Shoben (1983) demonstrate that sentence context can affect the time it takes to establish anaphoric reference between an exemplar (such as "football") and a category term (such as "ball"). Their data suggest that people entertain specific expectations about the nature of a vague term such as "ball" that are sensitive to the suitability of a particular exemplar in the context suggested by the sentence.

In answer to the question of how a listener adopts the appropriate understanding, Johnson-Laird (1981, 1983) has argued that it requires making inferences based on the *reference* of expressions rather than the sense or meaning. People's memory for sentences suggests that the product of the comprehension process is a model of the scenario depicted by a given sentence. Moreover, the model may contain details that are not specified by the meanings of the words in the actual sentence. For

example, Anderson & Ortony (1975) found "shark" served as a better recall cue than "fish" for (15).

- (15) A fish attacked the swimmer.

Similarly, Garnham (1979) has found that the word "fried" is a better recall cue than "cooked" for (16), but not for (17).

- (16) The housewife cooked the chips. [British English]
 (17) The housewife cooked the peas.

These data suggest that in processing a sentence such as (18), the word "shark" would serve as a better recall cue than "it."

- (18) It frightened the swimmer.

From this, however, we should not conclude that the listener must *always* consider "shark" as a possible meaning for "it." Johnson-Laird (1981) argues that comprehension of (18) requires the listener to imagine a scenario that is more specific than warranted by the meanings of the words. An adequate understanding of the utterance of a particular open-class word is inevitably underdetermined by an abstract characterization of its sense.

These sorts of observations imply that a sense-based explanation of language behavior will necessarily be incomplete. Whether or not a distinct *sense* is retrieved for lexical items, ultimately the language comprehension process involves the recruitment of knowledge about the likely nature of the particular referents of the utterances. Understanding "ball" in (12)–(14), for example, involves recruiting knowledge about the general scenario and the sort of ball that is most likely to be present in that scenario.

Consequently, a speaker who doesn't have the relevant cultural knowledge about baseball might entertain a different understanding of (14) than that of a sports-minded American. In contrast, an Indian speaker of English might be a bit puzzled by (14) and attempt to understand it as a strange variant on a cricket game. Regardless of whether there are shared truth functions between the Indian and the American speaker, there is a need to account for differences in their understanding. Given that the content and extent of a speaker's background knowledge affects her interpretation, it seems reasonable to postulate that integration with background knowledge is part of the interpretive process.

The dependence on a speaker's understanding of "ball" and her general conception of the scenario will become more obvious if we consider (19). "Ball" is understood in some less prototypical situation. For example, the understanding of "ball" in the context of the video game *NBA Jam*. Here, the "ball" is a basketball, but not a small sphere. Rather, it's a moving set of orange pixels, arranged in a roughly circular manner, covering the area of a few square centimeters. The ball in *NBA Jam* shares few physical features with a "real" basketball, but functions in a roughly analogous manner – that is, it is "dribbled" electronically and "passed" and "shot" by pushing buttons.

Moreover, the ball in *NBA Jam* is not the only two-dimensional entity that we call a "ball." We might point to a 4-inch picture of Shaquille O'Neal on the sports page and remark (as in (19)) on how small the "ball" looks in his hand.

- (19) That ball looks so small in his hand.

Or, we might point to a 40-foot billboard of Scottie Pippen and note the same thing.

Besides emphasizing the immense amount of contextual variability in the set of objects to which "ball" can be applied, these examples are also intended to highlight the systematicity in the disparate uses. For example, in both readings of (19), it is not the objective size of the "ball" that is being discussed, but the relative size of the ball and the man who is holding it. Similarly, in *NBA Jam*, there is a systematic set of correspondences between the objects and actions in the video game and the objects and actions in its three-dimensional counterpart. However, appreciating these correspondences requires the induction of an abstract schema.

Similarly, we can use the term "ball" to talk about the implications of televised sporting events. For example, during the Superbowl, Americans watching television might debate whether or not the "ball" was carried past the first-down marker. Alternatively, during the World Cup, those same Americans might scan the television screen in hopes of tracking the "ball." Just as knowledge of the game can help understand the significance of the game's events and aid the tracking of the ball, knowledge of the different games in which balls are used allows us to interpret statements about "balls" in ways that are far more specific than would be available from an abstract characterization of balls.

Now let's consider the case of two college students up late studying for an exam. Suddenly one crumples up a piece of paper and heaves it

at the wastepaper basket. The other decides to make a game of it and the two begin to shoot the "ball" at the "basket" from various spots around the room. After a while, one of them might utter (20).

(20) I'm getting good – I've gotten the ball into the basket the last three times.

Or after playing catch with a watermelon, a kid might utter (21) to explain what happened to dessert.

(21) We were playing catch till Jody dropped the ball.

Although neither of these cases involves a "real" game, we use the term "ball" to refer to the entity that functions similarly to a "real" ball in a "real" game.

We can also use the word "ball" to refer to entities that aren't directly involved in games. For example, imagine a bowling trophy with a small sculpture of a man, arm and leg extended, just about to release the "ball." The football coach might point to X's and O's on the chalkboard and say (22).

(22) By this point in the play, the ball is right here.

Recounting a playground basketball game with a much taller opponent, I might point to a spot above my head and say (23).

(23) So I'm going for the ball – which is way up here. . . .

Or, a musician might point to his saxophone and say (24).

(24) I'm Michael Jordan and this is my ball.

While philosophers have assumed a correspondence theory of truth in which the core cases of meaning involve a straightforward relationship between words and objects in the world, these examples highlight the extent to which language use often relates only obliquely to things in the world. It would be quite hard (probably impossible) to account for the occurrence of "ball" on the basis of objective criteria. Besides balls made out of leather and pigskin, we have balls made from fruit, trophy metal, and crumpled-up paper. We also have balls printed on newsprint, painted on wooden billboards, and composed entirely of moving pixels on the CRT. Some balls have physical extension, while others correspond only to a blank spot on the coach's chalkboard, or an empty spot in the air. Finally, it would appear that while some balls are for playing games, others are for melodies.

However, these sorts of arguments are bound to draw a predictable set of replies from the objectivist. For example, she could argue, quite rightly, that the term "ball" is merely vague. Surely, careful analysis will reveal an abstract common meaning to the referential, metaphoric, and represented balls. However, if this is the case, it's not clear what the abstract common meaning buys the language user. To capture the full range of cases, the meaning would have to be so abstract as to overlap with the meaning of other things, such as "saxophone." Further, it leaves unanswered the question of how the listener derives particular understandings of utterances of "ball."

Another retort is to suggest that the three-dimensional objects that are so much fun to throw, catch, and kick are the *core* cases while the others are parasitic extensions. However, this raises the issue of defining exactly which cases are to count as the core uses. Is the ball in *NBA Jam* a "core" instance of a ball or a representational extension? What about a toddler's toy football? My own intuition is that it depends on who's playing with the ball. If two toddlers are playing catch with it, it might count as a "core" instance of a ball. However, if two adults play with it, it becomes more "peripheral." Or what about a toy football used with a G.I. Joe doll? The problem is that there don't seem to be objective criteria that definitively determine the set of utterances that reflect the core.

One possibility is to let statistics adjudicate the core. On such a scheme, speakers will consider the way in which a term is used most frequently to indicate the core meaning. However, defining the core in this way does not help us to identify the essence of the concept. Moreover, there remains the problem of characterizing how the core meaning of "ball" can be extended. Although extension sometimes depends on physical resemblance to the core cases, there's no consistent dimension of similarity in which a "peripheral" type of ball is supposed to be measured. Sometimes color is relevant, sometimes it's not. Sometimes size is relevant, sometimes it's not. Sometimes relative size is relevant, sometimes it's not. As these examples are intended to demonstrate, an object can't be defined in terms of objective criteria. It's not the case, for example, that representational extensions of "ball" are all two-dimensional circular entities. Indeed, it would appear that the only property shared by all the balls discussed here is that they have been construed as balls.

The fact that so many instances of "ball" appear to be motivated by abstract relational features suggests that an adequate account of those relationships necessarily involves an appeal to human understanding. In particular, language use seems to reflect the establishment

of correspondences or *mappings* between elements and relations in different contexts. For example, representations such as the photograph, the billboard, or the television screen afford a systematic mapping of terms for elements in games to apply to splotches of ink, paint, and light. Although the entities in the representation are ontologically quite distinct from those in the actual game, visual similarities exist. Further, the representation contains information about the relationships between entities in the picture that correspond to relationships between the entities they depict, such as the relative size of a basketball and a basketball player's hand.

A crucial component of understanding any usage of "ball" involves knowing what portion of our vast knowledge about balls is to be attended to, and what should be ignored. In the previous examples, I have suggested that the meaning of a word in a particular context is intimately related to the nature of the context. Moreover, capturing the way in which this understanding varies with the context is impossible if we restrict ourselves to objective criteria of objects in the world. However, if we look instead to the cognitive processes that underlie meaning construction, we can understand the contextual variability of utterance meaning as reflecting differential access to information in memory as a function of context.

In the case of the watermelon, for example, its construal as a ball involves recruiting knowledge about how people typically play with balls, while at the same time ignoring what balls are typically made of. The case of the crumpled-up piece of paper is similar, and illustrates how the flexible use of "ball" can be extended to incorporate other entities in the college dorm room, such as the "basket" and indeed the students' "game" itself. Further, the use of terms in this way relies upon the constitutive nature of the participants' construal of their activity. The watermelon can be understood as a "ball" in (21) because the kids treat it *as* a ball; similarly, the students understand their antics *as* a game. This sort of language usage relies on the ability both to appreciate and to create cross-domain mappings.

Although people's behavior does not seem to reflect the use of a context-invariant sense, it does seem to reflect knowledge of typical scenarios in cases that provide no explicit context. For example, the difference between the understanding of the balls in (12)–(14) results from the way in which meaning inherently rests on an implicit set of assumptions about the world in general and (in this case) games in

particular. Further, when a proper context is provided, speakers can accommodate even atypical scenarios. In all these cases, the speaker exploits her knowledge about the context in order to create a scenario that contains the right sort of a ball. So while referential applications of "ball" appear to differ in kind from metaphoric and representational applications of the term, they merely reflect different portions of the spectrum of meaning construction.

1.2 BACKGROUND

While we tend to think of words like "cat," "mat," and "ball" as having simple definitions, the previous examples demonstrate the variability of meanings in concrete utterances – even in artificially constructed scenarios. Moreover, the particular character of utterance content seems to depend on contextual factors in some nonrandom ways. On this view, interdependence of meaning and background knowledge is no accident. Contextual variation in meaning is ubiquitous because context is an inherent component in the meaning construction process. Background knowledge affects utterance meaning precisely because utterances are designed to evoke information from memory in a way that supports action and interaction in the physical and social world.

This view stems from frame semantics and more generally from the framework of cognitive semantics (Fillmore, 1982; Lakoff, 1988). Rather than viewing the description of the world as the central function of semantics, the cognitive semanticist treats meaning as a cognitive phenomenon invoked to construe conceptual content. The cognitive semanticist is less interested in the truth conditions for sentences than she is in the cognitive operations that realize the production and comprehension of linguistic utterances. Because cognitive activity mediates the relationship between words and the world, the study of meaning is the study of how words arise in the context of human activity, and how they are used to evoke mental representations.

1.2.1 Frame Semantics

Fillmore (1982) suggests that the meaning of many words relies on speakers' experience with the scenarios and social institutions they presuppose. For example, words such as "Tuesday" cannot even be defined without providing a certain amount of background information about