1

Introduction

Figurative language generally refers to speech where speakers mean something other than what they literally say. Consider a situation in which Harriet says to a friend, "My marriage is an icebox." Most people recognize that Harriet intends to communicate something negative about her marriage, such as it is not emotionally affectionate or sexually passionate. How listeners arrive at these metaphorical meanings is a topic of considerable debate within both the humanities and cognitive science. One widely held belief is that metaphorical language must be more difficult to understand than plain, literal speech, precisely because metaphors, like all kinds of figurative language, are ornamental and deviate from literal thought. A standard way to assess this idea is to measure the amount of time it takes listeners, or readers, to comprehend figurative (in this case metaphorical) utterances compared to roughly equivalent literal speech.

The results of dozens of studies present a complex picture on this issue. Although, when seen in isolation, metaphorical utterances generally take longer to understand than literal ones, figurative speech can often be understood as quickly as literal speech when encountered in realistic discourse contexts (Gibbs, 1994, 2011 for reviews). This conclusion is especially true for more familiar, conventional figurative language, such as idioms (e.g., "kick the bucket"), stock metaphors (e.g., "John is a tiger"), conventional ironies (e.g., "A fine friend you are"), and certain indirect speech acts (e.g., "Can you pass the salt?"). Some scholars argue, nonetheless, that novel figurative language, such

as creative metaphors, will always require additional cognitive effort to understand compared to that needed to process nonfigurative speech. The extra time needed to comprehend figurative language is often assumed to reflect initial processing of a word's (or an entire phrase's) literal or semantic meaning, with context having a guiding role later in the interpretation process.

Our aim in this book is to critically evaluate the recent empirical work that examines people's successes and difficulties when interpreting figurative meanings. Although mainstream psycholinguistics did not initially pay much attention to the problem of how figurative speech was understood, because such language was seen as abnormal, there is now a huge literature that has examined many aspects of how people learn, produce, and comprehend figurative language. Not surprisingly, many theories have been proposed to account for how people understand different types of figurative language. We will describe these theories and the empirical studies testing them in the chapters that follow. Our primary focus is on the debates over the cognitive effort that is required to interpret speakers' and writers' meanings when they use figurative language, and the cognitive effects or meanings that arise from listeners' and readers' interpretive processes.

The time is ripe for this reappraisal of the research on figurative meaning because of some misguided assumptions about both the theories and empirical methods advocated in the study of figurative language processing. For instance, scholars continue to assume that there is a principled distinction between literal and nonliteral language, and that a single theory can account for how all aspects of figurative language are understood. We reject both of these beliefs.

At the same time, researchers, especially in cognitive science, maintain simplistic ideals about the nature of figurative meanings, such as the notion that speakers wish to convey single messages by their use of figurative, as opposed to so-called literal, speech. Scholars also assume that "understanding" figurative language rests on some

definitive "click of comprehension" that can be readily measured in scientific experiments. These idealistic views of figurative meaning have, in our view, brought forth a whole host of naïve theories that only address a small range of discourse situations in which figurative utterances are produced and interpreted.

We approach the topic of interpreting figurative meaning from the perspective of two cognitive psychologists who specialize in research exploring how, and why, people produce and understand language. Over the past 25 years or so, psycholinguistics has been primarily interested in describing the moment-by-moment mental processes that occur when people produce and comprehend language, ranging from the identifications of speech sounds and written words, up to larger meaningful units of conversation and texts. As psychologists have long realized, it is impossible to introspect about these rapid, mostly unconscious mental processes because humans have evolved to be remarkably efficient in using language. Although linguists, psychologists, and philosophers explore people's intuitions about aspects of language, such as judgments of grammaticality, cognitive scientists are unanimous in their belief that intuition alone is insufficient to systematically study the operations of the language processor. For this reason, psycholinguists (including neurolinguists) have developed a wide range of experimental methods that allow them to draw inferences about the automatic, unconscious processes occurring when people ordinarily use language. These methods are "indirect" in the sense that they require participants to engage in different tasks, the data from which can then be analyzed to draw conclusions about possible hypotheses on how figurative language is ordinarily understood.

But psycholinguists' primary interest with early aspects of processing leads them to posit general models of figurative language interpretation that might be quite misleading. People may not process all figurative meaning in the same way, precisely because the kinds and forms of different tropes are sufficiently varied as to resist

classification within a single theoretical framework. In a similar manner, psycholinguists must resist suggestions that there are different processing modes for literal and figurative language, or that parallel processes operate to produce literal and figurative meanings for utterances. Single linguistic processes may be sufficient to produce a variety of meanings without the need to postulate distinct processing modes for individual types of linguistic meaning.

We are strong advocates for conducting scientific research that tests potentially falsifiable hypotheses and, for this reason, believe that psycholinguistic methods provide an important set of tools for examining different theories on interpreting figurative meaning. The vast majority of the specific hypotheses examined by psycholinguists conducting figurative language research are rooted in ideas proposed by philosophers, linguists, and literary theorists. Translating ideas from other disciplines into a framework for conducting psycholinguistic experiments, however, poses significant challenges.

Many of the current debates within psycholinguistics on the cognitive effort involved in interpreting figurative meaning are rooted in deep misunderstandings of the literature outside experimental psychology. Scholars too often present simplistic views of different linguistic or philosophical theories, for example, which they then attempt to test using empirical methods that are inappropriate for the theories being investigated. In many cases, psychologists mistakenly assume that interdisciplinary theories of figurative meaning must reflect immediate comprehension processes, but there is no evidence that such an argument was ever intended by the linguists or philosophers who first proposed these ideas. More generally, there remains a significant gap between theory and empirical data within the psycholinguistic literature on interpreting figurative meaning. We draw attention to this problem, and suggest solutions to it in the chapters that follow. Furthermore, there are theories in linguistics, philosophy, and social psychology that have not received adequate attention by psycholinguists' experimental studies of figurative language

processing. We suggest, ways that these proposals can be readily tested within an experimental framework.

A central theme of this book is the uneasy tension between different disciplinary approaches to figurative language and how it is interpreted. Virtually all theories of figurative language processing within and outside experimental psycholinguistics have failed to address a major question: what meanings do people actually infer when a figurative word or expression is understood? For instance, how can one characterize what is understood when processing an expression such as "My marriage is an icebox"? The important emphasis on processing figurative language ignores exactly what people have understood when they seem to have successfully comprehended a particular figurative expression.

For the most part, psycholinguists and others tacitly assume that any figurative statement can be paraphrased by a similar linguistic expression that states in literal terms what people must have attempted to communicate when speaking figuratively (e.g., the American phrase "blow your stack" means "to get very angry"). This reduction of figurative meaning to simple, short linguistic paraphrases has in the past seemed reasonable in the context of designing experimental studies that, for instance, contrast figurative language processing with nonfigurative (or sometimes literal) discourse understanding. There is great effort put toward showing how people infer that something may have a metaphorical (as opposed to a non-metaphorical) meaning in some context, but we still do not know enough about the complex metaphorical meanings people create during rich pragmatic interpretation of figurative speech and writing. Psychological studies have shown a good deal about the social and emotional effects of irony understanding, for instance, but still do not offer a clear idea of what meanings arise during irony comprehension. Of course, the exact meanings people infer in discourse depend on many individual, linguistic, and situational factors. We need to further explore these different inferences in discourse with the long-term aim of showing

how different meaning products relate to different types of cognitive effort during figurative language production and understanding.

A concrete example of research related to this challenge investigated the pragmatic uses of metaphor (Gibbs & Tendahl, 2006). When people hear an expression such as "Lawyers are sharks" in discourse, their aim is not to simply understand the metaphoric meaning of this phrase, but to understand what pragmatic effect the speaker wishes to communicate by using this metaphor. For instance, in a conversation between two people, one person may state a number of negative thoughts about lawyers with a second person supporting this argument by saying "Lawyers are sharks." In this case, the metaphor simply strengthens the existing set of beliefs held by the conversational participants. But in a slightly different situation, one person may say several negative things about lawyers, to which the second speaker adds an additional negative assertion by uttering, "Lawyers are sharks." Finally, in a third situation, one speaker may comment on many positive attributes of lawyers to which the second person responds, "Lawyers are sharks," in order to contradict the first person. Thus, the same metaphor can achieve at least three different pragmatic effects (i.e., strengthening an existing idea, adding new information consistent with an existing idea, and contradicting an existing idea) depending on the context. Not surprisingly, people take more time to comprehend the metaphorical utterance "Lawyers are sharks" in the contradictory situation than in the other two.

One implication of this work is that understanding what any figurative utterance means is not simply a matter of getting to a particular figurative meaning, but understanding what a speaker pragmatically intends to achieve by use of that trope. Most reading time or brain scanning studies, for example, fail to consider these pragmatic effects by focusing exclusively on crude, simple distinctions between "literal" and "figurative" meanings, and incorrectly attribute variations in processing time or brain activity to constructions of those meanings as

opposed to the different social and pragmatic effects that speakers' utterances often convey in real discourse. In general, much greater attention is needed on the precise pragmatic effects achieved by different tropes and how particular figures of speech may lead to very different pragmatic and affective effects in varying contexts.

Furthermore, previous work on figurative language use has almost exclusively focused on people's understanding of single instances of figurative words or utterances. To date, research has not examined how people interpret very typical but nonetheless complex expressions which include numerous figures of speech such as seen in the following two sentences that opened a front page news story in the San Francisco Chronicle, titled "GOP in deep funk over Bush spending" (March 12, 2006):

> The Republican rebellion that President Bush smacked into with the Dubai ports deal was the tip of the iceberg of Republican discontent that is much deeper and more dangerous to the White House than a talk radio tempest over Arabs running U.S. ports.

People seem to understand *something* when they read sentences like this. How they do so, and what they actually interpret are two questions that should also be the focus of research, with the overarching goal of integrating such findings into related theories of sentence comprehension and discourse processing.

Some scholars suggest that there is a trade-off between the amount of cognitive effort put into linguistic understanding and the cognitive effects or meanings that are inferred (Sperber & Wilson, 1995). This theory applies to all aspects of linguistic communication, not just figurative language processing. Surely, in real time, people limit the effort put into understanding what others say. But we need to better understand both the factors that limit processing effort, and at the same time the actual effects (or meanings) that arise from understanding different types of figurative language. The problem of identifying, individuating, and perhaps counting, figurative

meanings is, therefore, a major issue for research on figurative language understanding. We also tackle this concern in the following chapters.

There is much at stake in research on figurative meaning as different theories of interpreting figurative meaning reflect contrasting conceptions of the human language processor, and, more generally, different versions of the relationship between thought and language. Concerns about the role of context in figurative language understanding mirror other continuing debates in psycholinguistics over the influence of nonlinguistic knowledge in language processing (e.g., studies on lexical ambiguity resolution and syntactic ambiguity in sentence processing). In this way, the research on interpreting figurative meaning is relevant to theoretical arguments on the modularity of language understanding and the figurative nature of human conceptual structures. As important as the larger theoretical debates are on the architecture of the language processor, and the recognition that studies on figurative language use can contribute important data to these debates, it seems evident that too many psycholinguistic studies use artificial stimuli constructed from researchers' own or students' raw intuitions. There is now an emerging literature in corpus linguistics that casts doubt on some of the stimuli used in figurative language experiments, enough so that scholars must be cautious when drawing theoretical conclusions about real-language use from the data obtained from many experimental studies. Corpus research provides guidelines by which better studies can be constructed to more closely approximate real-life figurative language use in context. We discuss this corpus work and its implications for studies interpreting figurative meaning.

We also argue that figurative language does not represent any cognitive or linguistic deviation, but in many instances directly reflects people's figurative conceptualizations of experience. One of us has forcefully presented this view in past writings (Gibbs, 1994, 2006a, 2011). Yet much new work has appeared in recent years, and the

time is right to again step back and survey the empirical/theoretical landscape on interpreting figurative meaning in relation to theories of human thought, and the embodied motivation for some forms of figurative language. Some hypotheses on interpreting figurative meaning can indeed be conclusively falsified. However, too many scholars advocate particular theories in the belief that their ideas are inconsistent with alternative views of how people comprehend figurative language. Psychologists, more specifically, often identify a single variable as important in figurative language processing, but then mistakenly postulate the existence of a single psychological mechanism to account for a particular empirical finding. This tendency to embrace the "empirical finding = psychological mechanism" fallacy has led to the creation of a vast number of different possible psychological processes that presumably are all operating whenever figurative meaning is understood. Once more, we reject this reflex to posit different psychological mechanisms for interpreting figurative meaning and suggest that a more comprehensive perspective may explain broader aspects of the cognitive effort and effects involved in figurative language use.

Our main goal is to construct a wide theoretical umbrella that can accommodate diverse empirical findings. One way to do this is to adopt an old, but still insightful, tetrahedral model of cognitive processes (Jenkins, 1979), which suggests that several broad factors shape processing, including (1) people (e.g., their abilities, interests, beliefs, motivations, goals), (2) the language materials (e.g., specific language, genre,), (3) the goal or task (e.g., understanding to solve a problem, make a decision, remember something, be emotionally affected by something said), and (4) the methods used to assess figurative performance (e.g., reading time, eye movements, brain scanning). These factors not only have their individual effects, but also interact in complex, even nonlinear, ways. In the end, we suggest that the complexities of figurative language processing are such that there may not be a single theory or model that specifically explains how all

aspects of figurative language are understood. In fact, it may even be a mistake, as suggested earlier, to believe that there is a unified type of language called "figurative language" that has its own special psychological processes which have evolved to enable people to efficiently process this kind of meaning.

Our alternative proposal is that the processes involved in understanding what is loosely referred to as figurative language are exactly those employed to interpret any instance of human speech. We argue for a broad theoretical framework, not specific to the study of figurative meaning, sometimes referred to as "dynamical systems models" or, more specifically, "constraint satisfaction models," that aim to capture the multiple, interacting influences on how people quickly comprehend, and on other occasions, more slowly interpret, various forms of figurative meaning in different contexts. These models, which are becoming increasingly popular in explaining many aspects of human performance in cognitive science, have the flexibility to, for example, describe how different aspects of contextual information may, or may not, facilitate the speed with which people interpret various figurative meanings, and what meanings people infer in context.

In this spirit, we hope to demonstrate how a dynamical systems approach to human cognition can be applied to account for the diversity of empirical findings on figurative language understanding, similar to that used to explain other facets of linguistic processing. At the same time, we maintain that many instances of figurative language convey special pragmatic effects that no other kind of speech can easily communicate. Another purpose of this book is to explore how the trade-off between minimizing cognitive effort and maximizing the special cognitive effects associated with figurative language can be scientifically studied and understood.

Chapter 2 addresses the question of what defines figurative meaning as opposed to any other kind of meaning or language. Psycholinguists typically assume that figurative meaning constitutes a conceptual category in which a speaker communicates something different