

Understanding and Literal Meaning

RAYMOND W. GIBBS, JR.,
University of California, Santa Cruz

The debate on literal meaning in theories of natural language processing has focused on two main issues: (1) Are there conditions by which the literal meaning of a sentence can be appropriately identified?; and (2) Is there some evaluation of literal meaning during the interpretation of natural language utterances? My article, "Literal meaning and psychological theory," (Gibbs, 1984) attempted to provide responses to these important questions from the perspective of contemporary psycholinguistic research. I argued that there is not a well-defined set of conditions for specifying the literal meaning of sentences in terms of compositional analysis and that the experimental evidence speaks negatively as to whether people *must* analyze the literal meanings of sentences as part of the process of understanding speakers' utterances. From these observations, I suggested that the widely accepted distinctions between literal and metaphoric meanings, and between semantics and pragmatics, have little utility for psychological theories of meaning and language use.

Dascal's (1987) recent defense of literal meaning raises some questions for my thesis. Although Dascal agrees with my essential point that literal meaning should not be equated with compositional meaning, he alleges that the concept of literal meaning can be profitably redefined as the conventional interpretation for a sentence. Contrary to my earlier position, Dascal states that with this revised definition of literal meaning one can easily see that some analysis of literal meaning plays a role, no matter how minor, in guiding understanders to the contextually appropriate interpretations of speakers' messages.

There are important implications for the resolution of the literal meaning debate. Virtually all theories of meaning in philosophy, linguistics, psychology, literary theory, and AI presuppose some, mostly implicit, view of what's literal and what's not. In this response to Dascal, I attempt to clarify

Preparation of this article was supported by Grant MH42980 from the National Institute of Mental Health and by a Faculty Research Grant from the University of California, Santa Cruz. I wish to thank Gayle Gonzales, Dom Massaro, Nancy McCarrell, Nandini Nayak, and Guy Van Orden for their advice on this paper.

Correspondence and requests for reprints should be sent to Raymond W. Gibbs, Jr., Program in Experimental Psychology, Clark Kerr Hall, University of California, Santa Cruz, CA 95064.

and extend my discussion of literal meaning and its purported role in understanding language. My position will be that literal meaning, whether viewed as compositional or as conventional meaning, is not an obligatory stage of analysis in psychological *process* models of language comprehension. People may phenomenologically identify some meanings as literal ones, but these are *products* of understanding and do not imply that different cognitive mechanisms drive the comprehension of literal and figurative discourse.

DEFINING LITERAL MEANING

Defining literal meaning is a difficult matter, one which surprisingly few theorists have attempted to do, despite the predominance of literal meaning in theories of understanding. Part of my 1984 paper was devoted to a critical discussion of the notion that the literal meaning of a sentence can be uniquely determined by the meanings of its component morphemes and by the syntactic rules of composition. It seemed clear that such an ideal was impossible to maintain given the variety of linguistic expressions with no well-defined literal meanings, but which can be readily understood.

Dascal (1987) agrees with this significant point. However, he proposes an alternative view of literal meaning, called "moderate literalism," which abandons the attempt to offer a set of necessary and sufficient conditions for something to be a literal meaning. In its place, Dascal suggests that several conditions and criteria are semantically relevant to literal meaning, but that no single condition or criterion is strictly necessary or sufficient. Even though compositionality is not sufficient to give a complete determination of literal meaning, it is still one source of information used in constructing the literal meaning of a sentence. This less formal description of literal meaning should allow one, Dascal argues, to include as part of the literal meanings of sentences, aspects of meaning such as hints, suggestions, and emotive meanings, along with the criteria of non-cancelability (in Grice's sense) and context invariance. Each of these aspects of meaning converges with compositionality to produce the literal, or as Dascal proposes, the conventional, meaning of a sentence. Presumably, this level of semantic representation plays an obligatory role in the process of leading listeners to identify the contextually appropriate meanings of utterances.

How are the various aspects of meaning and criteria, alluded to by Dascal, combined to produce a specific semantic representation that is *necessary* for understanding speaker meaning? It may be relatively easy to discern the conventional meaning of idiomatic phrases (e.g., *kick the bucket* literally or conventionally means "to die"). But it is not obvious how to apply the various criteria when determining the conventional meanings of utterances with innovative, figurative meanings. For example, is the literal meaning of Nietzsche's metaphorical assertion "Truth is a woman," or John Lennon's

lyric, "Happiness is a warm gun," related to these statements' compositional analyses or some other confluence of meanings more closely associated with their nonliteral interpretations? Dascal doesn't mention the likely possibility that composition and convention can give conflicting (i.e., incompatible) interpretations which only context can reconcile.

Although there may not be a set of necessary and sufficient conditions for specifying the literal meaning of a sentence, there should be some defined set of heuristics that are employed in determining the literal meaning of an utterance. I am not convinced that Dascal's list of meaning sources and criteria are themselves capable of generating literal meanings for sentences in a way that listeners can use such information when understanding what speakers mean. If Dascal is correct, the set of heuristics used for determining literal meaning should operate such that literal meaning can be automatically computed in real-time (i.e., within a few hundred milliseconds). However, Dascal fails to provide any reason to believe that his set of heuristics can actually specify unique literal meanings for each potential utterance in the language, given the constraints of on-line linguistic information processing.

At the same time, if the determination of literal meaning is a *necessary* part of understanding language, as Dascal assumes, then there must be some way of differentiating the criteria used in specifying literal meaning from those criteria, including literal meaning itself, that apply when understanding speaker meaning. But what are these differences? What sources of information are used in constructing the literal meaning for an utterance that are different from those used in comprehension of speaker meaning? In another paper, Dascal and Weizman (1987) propose a three-tiered organization of context and co-text to account for the various ways contextual and textual information operate in the determination of speaker meaning. However, this model still does not suggest a way of distinguishing between that set of information (i.e., background and context) which determines literal meaning from that set used to comprehend what speakers mean. Distinguishing between background and context won't help here because background knowledge is not part of meaning since it is nonrepresentational and preintentional (Searle, 1983). Furthermore, there is good reason to believe that there is no clear and distinct demarcation between intentional mental states and preintentional mental and bodily capacities (Johnson, 1987). Without some distinction, and details about *when* different information is used in constructing literal meaning as opposed to speaker meaning, there is no reason for viewing literal meaning as a special and obligatory part of understanding linguistic utterances.

The issue of *when* different knowledge sources are used in determining literal and speaker meanings is, of course, crucial to psycholinguistic accounts of language interpretation. Dascal's idea of *moderate literalism*

reveals a curious, and contradictory, view of linguistic information processing. On the one hand, Dascal (1987) asserts that he is "quite prepared to accept a model that postulates parallel or interactive processing" (p. 265). This admission is, in part, an acknowledgment of the difficulty in maintaining a strict separation between background and context in the understanding of utterances. But, on the other hand, Dascal still maintains that the rules of interpretation may "stipulate a preferred order of consultation of the various elements (literal meaning, different types and levels of context) that will yield an interpretation" (p. 265). Although this can "be in some cases overruled by other considerations" (p. 265), Dascal's advocacy of literal meaning necessitates that it be computed *before* the intended meaning of a speaker's utterance is determined. There is simply no way that Dascal can defend literal meaning without postulating a separate stage of analysis in which the literal representation of a sentence feeds into the process of recognizing the contextually appropriate meaning of an utterance. A truly parallel, interactive theory of language comprehension would not postulate that literal meaning is necessarily determined in some autonomous stage of processing apart from the influence of other (contextual) sources of information used to interpret speaker meaning.

The problem here is that the aspects of meaning and criteria, supposedly used in determining literal meaning, do not converge in any natural way. It appears that some of the sources of information contributing to literal meaning may themselves be *products* of interpretive acts. Recognizing that an utterance is a hint, or has some particular emotive meaning is to have already understood, at least partially, what a speaker intends. If a listener uses different information as part of the criteria for determining a sentence's literal meaning, which itself requires an interpretive act, why can't that same information be used in the actual interpretation of the speaker's utterance? One cannot argue that recognition of literal meaning is an autonomous stage of linguistic processing if the determination of literal meaning is influenced by other contextual factors used in the comprehension of speaker meaning.

Although Dascal maintains the traditional assumption that literal meaning is partially defined as context-invariance meaning, what is seen as a hint or emotive meaning requires that some understanding of the utterance in *context* be made. Even the determination of conventional meaning is context-driven (Gibbs, 1986a). Speakers might agree that the conventional meaning of the indirect speech act, *Can you pass the salt?* is "Pass the salt." But this meaning presupposes a typical context of use where salt is being requested and not where one's ability to pass the salt is questioned. The conventional interpretation of an utterance presupposes some context of use perhaps so widely shared that there *seems* to be no context at all.

People often fail to acknowledge that certain knowledge and assumptions drive their judgments of literalness. For instance, listeners will state that the

expression, *The cat is on the mat*, "literally" means that the cat is on top of the mat even though this judgment is an interpretative act making use of widely shared assumptions of some specific contextual framework (e.g., the existence of gravity, see Searle, 1978). It is virtually impossible to make judgments about the literal or conventional meaning of an utterance apart from some actual or supposed context. The negative consequences of judging literalness apart from some communicative situation is that we end up confusing the *judgment* of meaning with the *process* of comprehending meaning.

LITERAL MEANING AND THE PROCESS OF UNDERSTANDING

Dascal claims that the alleged empirical evidence against literal meaning, reviewed in my 1984 paper, actually supports *his* view of literal meaning and its purported role in language processing. These experimental findings generally indicate that people take no more, and in many cases, less time, to process the nonliteral interpretations of idioms, indirect speech acts, sarcasm, and metaphor than to comprehend literal interpretations. Dascal assumes that because these experiments primarily examined comprehension of conventionalized figurative speech, the data really lead to the conclusion that the conventional meanings of utterances are, in fact, their literal ones. Dascal goes on to challenge: [To] "convincingly show that literal meanings are unnecessary one would have to devise experiments which prove that they play no role in the interpretation of such clear-cut cases of indirection as conversational implicatures, radically novel metaphors, and so on" (p. 267).

My claim is that such data already exist. Many psycholinguistic experiments have shown that novel metaphorical utterances, such as, *Billboards are warts on the landscape*, can take no longer to understand than literal uses of the same sentences (Ortony, Shallert, Reynolds, & Antos, 1978; Shinjo & Myers, 1987), or nonmetaphorical equivalent statements (Inhoff, Lima, & Carroll, 1984). There are also data on understanding of sarcasm and irony which indicate that people often take *less* time to understand the conveyed, sarcastic meanings of comments like, *He's the best driver I've ever seen* (meaning "He's a terrible driver"), or sarcastic indirect requests, such as, *Why don't you take you time getting the ball?* meaning "Hurry up and get the ball," than to comprehend literal uses of the same sentences or nonsarcastic equivalent sentences (Gibbs, 1986b, 1986c). All of these studies clearly show that the interpretation of novel figurative language can occur without a preliminary analysis of these sentences' literal meanings.

Dascal dismisses the relevance of the Gibbs (1986c) data by conjecturing that the sarcastic, indirect meanings of these utterances may have been "particularly perspicuous or salient in the context" (p. 268). But the very

point of these psycholinguistic experiments was to investigate the influence of contextual information on understanding figurative language. With appropriate context, people need not analyze the literal (or compositional) meaning of a figurative utterance before deriving its nonliteral interpretation. Dascal's argument that the data on sarcastic indirect requests are irrelevant to the literal-meaning debate adds up to a dismissal of the context-situated reality of everyday language processing.

It is not surprising that Dascal refuses to give much weight to the empirical evidence in my 1984 paper. He contends that although literal meaning must play some role in understanding the contextually appropriate meaning of a speaker's utterance, "the rather subtle nature of this control or guiding role may well be responsible for some of the empirical difficulties in detecting the presence of literal meaning in certain experimental studies" (p. 262). Dascal wants us to accept, almost as a matter of faith, that literal meaning exists and plays an obligatory role in comprehending speech, despite the empirical evidence to the contrary. But if the guiding role that literal meaning has in linguistic processing is potentially so subtle as to be undetectable, then Dascal has essentially asked us to embrace a theory that is nearly impossible to falsify. The burden of proof, however, lies with the proponents of the *literal meaning hypothesis*, described in my 1984 paper, or Dascal's view of *moderate literalism*, to find evidence that literal meaning *does* play a guiding role in the interpretation of utterances. This burden extends to theorists who hold that literal meaning is computed in parallel with other, perhaps more contextually salient, meanings.

LITERAL MEANING AND THE PRODUCTS OF UNDERSTANDING

One of the prime motivations for maintaining a belief in literal meaning stems from the phenomenological experience that different kinds of meanings are recovered during the interpretation of metaphoric and literal language. The indeterminate nature of metaphor makes its meanings seem "special," and by opposition, literal meanings seem ordinary and primary. Where does the intuition that metaphor seems "special" and different from literal language arise from? Many defenders of literal meaning appeal to their phenomenological experience that metaphoric meaning somehow feels different, as motivation for the idea that different cognitive mechanisms drive the comprehension of literal and metaphoric speech.

Dascal demonstrates a perfect example of this tendency in his discussion of jokes, hypnosis, and dreams. When understanding a joke, there seems to be a moment when we quickly recover a different meaning from a line or phrase that we previously believed to have correctly understood. But why must one of a joke's potential meanings, through which the joke achieves its humor, be the "literal" meaning, as Dascal argues, while the other is referred

to as the joke's "secondary" or "metaphorical" meaning? The belief in a literal or primary meaning assumes that the humor in jokes lies in the text itself rather than in the (mis)interpretations of the speakers' intentions. Thus, Freud's joke where one man asks, "Have you taken a bath?" and the second man replies asking the other in return, "Why? Is there one missing?" seems funny precisely because the listener misinterprets the speaker's actual intention in stating the question, *Have you taken a bath?* in the later context of a missing bath. It is the shifting of speaker's *intentions* through which a joke gets its humor, not in the shift from a "literal" to "secondary" meaning in the text. It makes no sense to call one possible intention a "literal meaning" unless one can state exactly what makes this meaning psychologically primary. All Dascal has done, is to show, quite correctly, that there are a variety of *products* that result from language understanding. Nobody disputes this as an aspect of psychological reality. I do, however, dispute the notion that some of these products are primary and reflect the output of a unique cognitive process.

A similar confusion between the *process* and *products* of comprehension can be found in Dascal's analysis of hypnosis and dreams. A hypnotized subject's ability to imagine the sensation of not having hands, although still knowing that the hands exist, only reflects an awareness of different possible realities, not different, and cognitively separable, literal and nonliteral *meanings*. Even when a person has a dream depicting a bare chest and interprets it as expressing the idea of "getting something off one's chest," this only reflects different *products* of interpretation. It does not necessarily imply that one interpretation (i.e., the literal) is more primary or fundamental to human experience than the other (i.e., the nonliteral).

People can sometimes judge some statements as literal and others as metaphorical. But do these judgments necessarily reflect different cognitive mechanisms underlying comprehension of literal and metaphorical language? I suggest that metaphor appears "special" precisely because of the *products* of comprehension, and not because of the *process* by which metaphorical meanings are understood (Gibbs & Gerrig, in press). Metaphor is often intended to convey various indeterminate, nonpropositional meanings whose recovery is highly constrained by the context of shared beliefs held by speakers and listeners. And it is the exploitations of this context that make the use of metaphor appear to be a "special" psychological activity.

Consider the following brief conversational exchange (from Gibbs & Gerrig, in press).

Rob: Does Gladys have a good memory?

Denise: Gladys is just like an elephant.

Why does Denise respond to Rob's question with a metaphorical assertion, and not a simple, literal statement? We assume that Denise wants Rob to use his encyclopedic knowledge of elephants, including, let's suppose, the

folklore that elephants are reputed to have excellent memories. With this information, Rob could infer that "Gladys has a very good memory." Although Denise expects Rob to recover the implicature that "Gladys has a very good memory," and all the implications from it that Rob might be interested in (i.e., regarding Gladys's phenomenal memory powers), it would be unusual if this implicature is *all* Denise expected Rob to infer. After all, Denise could have explicitly stated, *Gladys has a very good memory*, in response to Rob's question.

But metaphorical talk often presupposes and reinforces an intimacy between speaker and listener, the cultivation of which is, perhaps, the primary function of such language (Cohen, 1979; Gibbs, 1987). Intimacy can be enjoyed by all those who are confident that what they say will be understood. It is the bond among those who share not only a basic linguistic competence, but a common stock of experiences, interests, sensibilities, and the ability to call upon that information when interpreting language. In the above example, Denise may want to cultivate some degree of intimacy between herself and Rob by her use of a metaphorical description of Gladys's memory. Suppose that Gladys *really* is quite large and could, under other circumstances, be described as being as large as an elephant. If this belief were part of Denise and Rob's "common ground," then responding to Rob's question by saying, *Gladys is just like an elephant*, would be ironic, in that Denise specifically intends for Rob to infer both "Gladys has a very good memory," and "Gladys is quite large" (but not the unauthorized inference that "Gladys has a large nose."). It is not unreasonable to assume that Denise could intend Rob to draw these inferences. If Rob failed to recover Denise's intention to achieve a mutuality, a sense of intimacy or complicity, then it might be incorrect to assume that Rob fully understands the meanings of the metaphor.

It is recovery of these nonpropositional, sometimes poetic, effects that makes metaphor seem so different from so-called literal language. Many other kinds of nonliteral language seem special in much the same way. For example, when a speaker makes a sarcastic remark, such as, *This is really a delicious dinner*, he intends to convey not only some propositional information (e.g., "This is a terrible dinner"), but something about his attitude or opinion about this proposition. Indeed, it is precisely because they are evocative that we have conventionalized, idiomatic expressions; such as *lick one's wounds*, *knock one's socks off*, or *steal one's thunder*. Metonymic statements, such as Nunberg's (1979) example of one waitress saying to another, "The ham sandwich is getting impatient for his check," create similar affective import. Note that it is immaterial whether listeners need actually *calculate* these inferences each time they hear either conventional or truly poetic metaphors. All that matters is that people recognize that the use of a figurative phrase is motivated by a desire to communicate various nonpropositional, aesthetic meanings.

What makes metaphor appear "special," and so different from literal speech, are the *products* of this "common ground" interplay, the results of the comprehension processes as constrained by context, not the *process* itself. Dascal, and other defenders of literal meaning's role in figurative language processing, mistakenly confuse the *products* of comprehension with the psychological *processes* used to achieve those products. There are occasions when we perceive some piece of language as being "literal" or "metaphorical," but this in *no way* proves that distinct cognitive mechanisms *must* be needed to arrive at these meanings.

■ Original Submission Date: April 22, 1988.

REFERENCES

- Cohen, T. (1979). Metaphor and the cultivation of intimacy. In S. Sacks (Ed.), *On metaphor*. Chicago: University of Chicago Press.
- Dascal, M. (1987). Defending literal meaning. *Cognitive Science*, 11, 259-281.
- Dascal, M., & Weizman, E. (1987). Contextual exploitation of interpretation in text understanding: An integrated model. In M. Papi & J. Verschueren (Eds.), *The pragmatic perspective*. Amsterdam, Netherlands: John Benjamins.
- Gibbs, R. (1984). Literal meaning and psychological theory. *Cognitive Science*, 8, 275-304.
- Gibbs, R. (1986a). What makes some indirect speech acts conventional? *Journal of Memory and Language*, 25, 181-196.
- Gibbs, R. (1986b). On the psycholinguistics of sarcasm. *Journal of Experimental Psychology: General*, 115, 1-13.
- Gibbs, R. (1986c). Comprehension and memory for nonliteral utterances: The problem of sarcastic indirect requests. *Acta Psychologica*, 62, 41-57.
- Gibbs, R. (1987). What does it mean for a metaphor to be understood? In R. Haskell (Ed.), *Cognition and symbolic structures: The psychology of metaphoric transformation*. Norwood, NJ: Ablex.
- Gibbs, R., & Gerrig, R. (in press). How context makes metaphor comprehension seem "special". *Metaphor and Symbolic Activity*.
- Inhoff, A., Lima, S., & Carroll, P. (1984). Contextual effects on metaphor comprehension in reading. *Memory and Cognition*, 12, 558-567.
- Johnson, M. (1987). *The body in the mind*. Chicago: University of Chicago Press.
- Nunberg, G. (1979). The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy*, 3, 143-184.
- Ortony, A., Schallert, D., Reynolds, R., & Antos, S. (1978). Interpreting metaphors and idioms: Some effects of context on comprehension. *Journal of Verbal Learning and Verbal Behavior*, 17, 465-477.
- Searle, J. (1978). Literal meaning. *Erkenntnis*, 13, 207-224.
- Searle, J. (1983). *Intentionality*. Cambridge, England: Cambridge University Press.
- Shinjo, M., & Myers, J. (1987). The role of context in metaphor comprehension. *Journal of Memory and Language*, 26, 226-241.