

## A Framework for a Theory of Meaning

Robert Wilensky  
Division of Computer Science  
Department of EECS  
University of California, Berkeley

The notion of meaning is central to theories of language. However, there appears to be considerable confusion regarding what a theory of meaning should do, and how it pertains to other linguistic issues. In this paper, I attempt to rectify one aspect of this confusion, namely, the relation of literal meaning, sentence meaning, and speaker meaning. To do so, I present two polar opposites, a standard orthodoxy and a radical challenge. I then attempt to show that, as in most cases, the truth is not in between, but rather, requires a completely different framework.

### The Right-Wing Orthodoxy

First, there is a widely held orthodoxy, which I refer to as the “right wing” position on meaning. The right wing meaning dogma includes the following assumptions:

(1) There is a meaning that can be associated with a given sentence independent of context, usage, or speaker. This is known as the *sentence meaning*. Of course, a given sentence may have several such meanings, in which case it is ambiguous, or none, in which case it is anomalous.

(a) Sentence meaning is *compositional*. That is, it may be determined from word meanings together with general rules of grammatical construction.

(b) The sentence meaning is the same as the *literal meaning* of a sentence. While there may be many different interpretations that we can assign to a sentence, these do not belong to the sentence per se. Rather, they require recourse to the context, to the speaker’s intentions, to extra-grammatical knowledge, and so on. Only the literal/sentence meaning is strictly a function of the sentence alone.

(c) The sentence meaning (of a sentence in the indicative) establishes a set of truth conditions for that sentence. Most truth-theoretical accounts equate these truth conditions with the meaning of the sentence.

(d) The sentence meaning is completely independent of context. While the truth of a sentence’s meaning may vary with context, and other interpretations may become available in different contexts, the sentence meaning remains constant over all such variation. Thus, while the truth of the sentence “Today is Tuesday” is a function of when it is uttered, and that of “My name is Peter Smith” a function of who does the uttering, the literal meaning, i. e., the set of conditions that determines whether the sentences are true, is for both sentences invariant under such changes.

In practice, a stronger form of independence is presumed. This is that sentence meaning is that interpretation of a sentence that can be made in the “zero” or “null” context. I call this ultra-right wing assumption the *strong literal meaning position*, as it commits one to equating a literal meaning with a particular interpretation of an utterance.

(2) The other possible interpretations we may wish to assign to a sentence are distinct from the sentence meaning itself, and are usually classified as part of the *speaker’s meaning*. These interpretations may differ radically from the sentence meaning. For example, in the case of sarcasm, the speaker’s meaning might be the opposite of the sentence meaning. The speaker might mean something quite apart from the sentence meaning if he is intending the sentence to be interpreted idiomatically, or if he is using any one of a repertoire of linguistic devices, such as metaphor, metonymy or synecdoche. The speaker may mean something related to but different from sentence meaning if the sentence is an instance of exaggeration or meiosis. In

most cases, even if the sentence is meant to be interpreted “literally” the speaker means something in addition to the sentence meaning, by way of conversational implication, conventional implication, the use of indirect speech acts, and the like.

(a) To say that the speaker’s meaning of a sentence is different from the sentence’s meaning does not deny that such meanings are determined by cultural or linguistic conventions. For example, an idiomatic expression like “kick the bucket” requires a linguistic convention for it to have its intended effect, while finding the speaker’s intention underlying an indirect speech act like “Do you know what time it is?” may presume a cultural one. The various interpretations of the speaker’s meaning may have conventional underpinnings, but these are apart from those that assign meanings to sentences of the language.

(b) In some cases, the speaker intends to communicate exactly and only the sentence meaning of a given utterance. In such cases, the speaker might be described as saying exactly what he means.

Note that this is a special case of the strong literal meaning position, in which (i) one interpretation of the sentence is its literal meaning, and (ii) the speaker wishes to communicate exactly this content.

(c) Literal meanings form the basis for determining the speaker meaning of a sentence. According to some accounts, a hearer trying to interpret a sentence first attempts to determine that sentence’s literal meaning; then this meaning is judged to be deficient in some way, and another (speaker) meaning is determined in its place. Thus, a speaker hearing the sentence “My car guzzles gasoline” finds this sentence meaning anomalous, and then determines that the speaker must have desired that the sentence be interpreted metaphorically. Similarly, a hearer trying to interpret the sentence “That was a brilliant idea” in a context in which he had done an obviously stupid thing first determines the literal or sentence meaning of this utterance. This is then judged as inappropriate under the circumstances, so the hearer eventually infers that a sarcastic interpretation is warranted.

It is important to emphasize that, according to the orthodox position, all the various interpretations we may wish to associate with a sentence other than its literal meaning are something other than the sentence meaning itself. Rather, they are attributable to the *use* of that sentence by a given speaker in a given situation. Thus, the sentence meaning of “That was a brilliant idea” is always the literal interpretation of that sentence, namely, that some particular idea was clever. It makes no difference if this interpretation is clearly ruled out by the situation, and could not possibly have been what the speaker intended.

### The Left-Wing Challenge

This right wing orthodoxy has been subject to radical challenges from the intellectual “far left”. We might term such challenges *contextualist* theories, as their advocates believe that it is not possible to talk about meanings of sentences apart from the particular meanings that such sentences take on in particular contexts. Thus the notion of a sentence meaning is rejected outright by contextualists.

Contextualist positions arise in a number of quarters. For example, consider the attempts to defend Austinian “use theories” of meaning, that is, theories which attempt to explain the meanings of sentences in terms of the uses of sentences. One notable attempt is that of Searle (1979). Searle’s begins by considering sentences which are commonly believed to have clear literal meanings, such as

(1) The cat is on the mat.

Searle then argues that it is not possible to assign this sentence a literal meaning independent of context, because there are a number of (admittedly bizarre) contexts in which the literal meaning does not seem to determine the sentence’s truth conditions. For example, Searle considers the case in which the cat is, unbeknownst to its owner, drugged into stiffness and balanced on the edge of a mat which is itself firm, and at an angle to the ground. In this case, if I respond to the owner’s question about the location of the mat by

saying (1), Searle claims that this utterance “should probably be described as an ingenious lie”. However, suppose the same sentence is used in a situation in which both the speaker and hearer are aware of the unusual orientation of the mat, and in which nearby there are other objects the cat might be on. In this context, uttering (1) seems appropriate. Searle suggests that this example casts doubt on the thesis that there exists a context-independent literal meaning that determines the truth conditions of that utterance in every context.

In a similar vein, Searle (1980) argues that the semantic content of ‘cut’ makes different contributions to sentences like the following:

- (2) Cut the grass
- (3) Cut the cake

Searle claims that ‘cut’ makes a different contribution in each case because, if it turned out that I ran the cake over with a lawn mower, say, I would not have complied with the speaker’s intention.

In Searle’s analysis, we must abandon an important tenet of the orthodox position, namely, that the literal meaning of a sentence determines the truth conditions of that sentence.

### The Right-Wing Defense

Katz (1981) disposes of this sort of objection as follows. He points out that Searle confuses sentence meaning with sentence use in making his claims. That is, uttering (1) in the first situation may indeed deceive the cat’s owner. But deception involves a theory of sentence use, not merely one of sentence meaning. Similarly, we may be conveying precisely what is understood by uttering (1) in the latter situation, but this too is a matter of sentence use. The fact that the (presumably identical) sentence meaning can contribute to deception in one circumstance and convey one’s intentions accurately in another should not bother advocates of the orthodox position any more than the fact that the same utterance might be taken literally in one context and ironically in another. In both cases, the sentence meaning is same. But its contribution to the communicated meaning is completely different.

Katz raises the same objection to Searle’s claim that the semantic content of ‘cut’ makes different contributions to sentences (2) and (3). Agreeing with Searle that a listener who ran over the cake with a lawn mower would not have complied with the speaker’s intention, Katz observes that the *speaker’s intentions* are not the issue. The issue is the sentence meaning, and there is no reason that sentence meaning alone should have to determine compliance, etc., with a speech act.

In sum, according to Katz, the orthodox position does not assume that the literal meaning of a sentence *determines* the conditions of satisfaction of the *use* of that sentence as an utterance. It merely contributes to such conditions, perhaps in an involved and open-ended way. However, these complexities are complexities of sentence use, not of sentence meaning. Therefore, sentences such as (1), (2) and (3) are not troublesome. Rather, their literal meaning participates in the determination of the meaning they would convey in an actual situation. Given this larger view, there is no reason that we should, as Searle suggests, give up on the idea that literal meaning determines the truth conditions of a sentence; Searle has merely shown that the meaning of a sentence is only one of the factors entering into the meaning of an utterance. But the latter involves a theory of sentence use, not of sentences per se, so the orthodox position is quite safe.

Katz’s counters to Searle’s objections as stated seem correct. In particular, most of the force of Searle’s argument is that utterance meaning typically departs from sentence meaning, and it is utterance meaning that determines whether a request is satisfied or an order obeyed, etc. Since the actual contribution of sentence meaning to utterance meaning is left unspecified by the orthodox position, the objection cannot harm it.

## The New Problem

However, in making his point, Katz's opens the door for an even more serious objection. This is as follows: To accept Katz's position, we must allow that sentences like (1-3) have a meaning that is different from that we would assign the uses of these sentences in the null context. Consider the case of "The cat is on the mat". Now, the preposition "on" can be used to express quite a few different physical situations. If there is a central meaning to "on" that encompasses all of these, it would seem to mean something closer to "supported by" than "physically above and in contact with", as both Searle and Katz seem to presume. That is, in addition to meaning "lying on" and "balancing on", we use "on" literally to mean "hanging from" (as in "the fixture on the ceiling") and "vertically supporting" (as in "the notice on the bulletin board"). While the dictionary definition of "on" would appear to be neutral with respect to a great number of physical relationships, it appears as if "on" can only be used to mean one of these more specific relations. That is, when we say "The cat is on the mat", it is hard to imagine that an utterance of this sentence could ever remain neutral with respect to whether the cat is lying on the mat, attached to the mat a la the bulletin board, or suspended from the ceiling in the manner of the light fixture.

An important tenet of the right wing orthodoxy has been undermined. This is the "strong literal meaning" position, in which the sentence meaning can, in some limiting case, serve as utterance meaning. According to our argument, sentences like "The cat is on the mat" have a perfectly well defined "sentence meaning" – it's just that one can never *mean* it.

Thus it seems that "sentence meaning" is not really a meaning after all. That is, that object that one can compute using the grammar and lexicon may very well *never* be in itself a suitable candidate for the meaning of an utterance. This object may be related to actual meanings in important ways; but that should not confer meaningfulness upon it. Indeed, one of the basic and erroneous tenets of the right wing position seems to be predicated on this terminology alone: While there may indeed be a semantic object computable from the grammar and lexicon without recourse to context, the interpretation of this sentence when no external context is supplied is likely to depart from this object. Our knowledge of the world suggests to us that "Cut the cake" refers to slicing, even if no further context is supplied. So the interpretation of a utterance in the null context may be different from its sentence meaning.

Moreover, it seems that the term "literal meaning" is inappropriate throughout all these discussions. It seems that the term "literal" is useful only for talking about interpretations. Thus, it doesn't make any sense to talk about literal meaning of a *sentence* even an unambiguous one. The cat being balanced on the edge of a mat is a literal *interpretation* of the sentence "The cat is on the mat", and the cat lying on the mat is a literal *interpretation* of "The cat is on the mat". But neither can be distinguished as a privileged literal meaning of a sentence.

In sum, there may be an object that can be computed from the grammar and the lexicon. However, such an object may not in itself be a suitable candidate for a meaning. Moreover, the "literal meaning" of a sentence is something other than the sentence meaning, in that it seems to include interpretations of the sentence based on default knowledge, even if no additional context is supplied. So "sentence meaning" may not be a meaning at all, and "literal meaning" is something radically different from sentence meaning.

Once literal meaning is divorced from sentence, then there appears to be no motivation for brutally distinguishing "core" grammatical considerations from less generative linguistic conventions. That is, we can include all linguistic constructions in the domain of grammatical knowledge when talking about computing something from a grammar and lexicon. In this mode, sentences such as

(6) John kicked the bucket.

are *grammatically* ambiguous, rather than having grammatical sentence meaning and an extragrammatical speaker meaning. Thus, we can associate the idiomatic interpretation with the sentence itself, rather than uncomfortably fitting this fact into the domain of speaker meaning.

## The New Picture

These challenges to the orthodox right wing position seem to dictate the following new picture: There is a semantic object we can assign to a sentence independent of context, based on a grammar and lexicon; however our notion of grammar is broad enough to include non-core constructions. While this object is a semantic object may not in itself be a suitable candidate for a meaning.

This framework requires a new terminology. We will call the semantic object assignable to a sentence the sentence's *primal content*. This object generally will be interpreted to become a meaning, with its content restricted by world knowledge, etc. This conveyed meaning is called the *actual content* of the utterance.

For example, in the case of "The cat is on the mat", the primal content will be the rather abstract proposition that the cat is somehow supported by the mat. The actual content, based on pragmatic knowledge of cats, mats, etc., is some "lying upon" relation. Similarly, the primal content of "Cut the grass" and "Cut the cake" are identical insofar as the prescribed action is concerned, while the actual content of the first would most likely refer to mowing, and of the latter, to slicing. We have referred to this particular inference type, namely, that of inferring a more precise interpretation from a vaguer semantic content, as *concretion* (Wilensky 1983, Norvig 1986).

Furthermore, since non-productive constructions are included in the grammar, the primal content of some sentence can include that which is beyond sentence meaning in the orthodox dogma. For example, the primal content of "John kicked the bucket" could refer to dying as well as kicking. In addition, analyses that make recourse to conventionalized metaphor (e. g., Jacobs 1985, Lakoff and Johnson 1980) fit in nicely within this framework. Thus a sentence like

(7) John gave Mary a kiss.

would be said to have a primal content that involves transferring, and an actual content that involves kissing. The process of using conventionalized metaphorical knowledge (called *views* in Wilensky (1986)) to interpret a primal content is called *unviewing*.

In sum, there are a number of dichotomies in the orthodox position that are worth preserving. The problem is that they have been bundled together erroneously. I offer the primal/actual content distinction as a more helpful framework to talk about the meaning of sentences and utterances.

## References

- (1) Jacobs, P. S. A knowledge-based approach to language production. Report no. UCB/CSD 86/254. Computer Science Division (EECS), UC Berkeley. August, 1985.
- (2) Katz, J. Literal meaning and logical theory. *Journal of Philosophy*, 78 pp. 203-234. 1981.
- (3) Lakoff, G. and Johnson, M. *Metaphors We Live By*. University of Chicago Press, Chicago, Ill. 1980.
- (4) Searle, J. Literal meaning. In J. Searle (ed.). *Expression and Meaning*. Cambridge, England, Cambridge University Press. 1979.
- (5) Searle, J. The background of meaning. In J. Searle, F. Keifer, & M. Bierwisch (eds.), *Speech Act Theory and Pragmatics*. Boston, Mass., D. Reidel. 1980.

- (6) Wilensky, Robert. *Planning and Understanding: A Computational Approach to Human Reasoning*. Addison-Wesley, Reading, Mass. 1983.
- (7) Wilensky, Robert. Some Problems and Proposals for Knowledge Representation. Berkeley Computer Science Technical Report UCB/CSD 86/294, May 1986.