

Primal Content and Actual Content:
An Antidote to Literal Meaning*

Running Head: An Antidote to Literal Meaning

Robert Wilensky

Division of Computer Science
Department of EECS
University of California, Berkeley
Berkeley, CA. 94720

*This research was sponsored in part by the Defense Advanced Research Projects Agency (DOD), ARPA order No. 4871, monitored by Space and Naval Warfare Systems Command under Contract N00039-84-C-0089, and by the Office of Naval Research under contract N00014-80-C-0732.

Abstract

It has generally been assumed that, since the term "literal" distinguishes productive uses of words from idiomatic uses, non-metaphoric from metaphoric, and direct from indirect, that literal meanings must be the same as sentence meanings, i. e., that they could be computed from knowledge of the words and core grammar rules of the language.

However, this widespread presupposition appears to be false. In particular, literal interpretations of a sentence, even out of context, generally make recourse to extra-linguistic knowledge, while some non-literal interpretations are purely linguistic in nature. Furthermore, the semantic content derivable from purely grammatical and lexical knowledge may not even be a possible interpretation of a sentence.

Since these distinctions are hopelessly misleading, a new set is proposed based on a very different organization of knowledge. "Primal content" refers to the interpretation we can assign to a sentence based on lexical and grammatical knowledge, broadly construed. "Actual content" refers to the specific meanings speakers encode into utterances and extract out of utterances, generally making liberal use of extra-linguistic facts. The resulting dichotomy is meant to provide a firmer basis for theorizing about meaning.

1. Introduction

The notion of meaning is central to theories of language. However, there appears to be considerable disagreement regarding what a theory of meaning should do, and how it pertains to other linguistic issues. While these assumptions about meaning are rarely made explicit, they are nevertheless crucial, as numerous controversies in linguistic theory appear to be attributable more to terminological abuse than to more substantive theoretical issues.

In this paper, I attempt to clarify one aspect of this issue, namely, the relation of literal meaning, sentence meaning, and speaker meaning. To do so, I first present two polar opposites, a standard orthodoxy and some radical challenges. I then attempt to show that, while the orthodox position survives the challenges intact, it is flawed in serious ways. To salvage the many important insights that the orthodoxy has to offer, but to avoid its difficulties, a new dichotomy, called the *primallactual content* distinction, is considered.

2. The Right-Wing Orthodoxy

First, let us examine what I view as 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) Compositionality - Sentence meaning is *compositional*. That is, it may be determined from word meanings together with general rules of grammatical construction.

(b) The Equivalence of Sentence and Literal Meaning - 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 or sentence

meaning is strictly a function of the sentence alone.

- (c) Truth Conditionality - 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.

Note that, according to (a), 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 as follows:

- (d) Null Context - The sentence meaning is that interpretation of a sentence that can be made in the "zero" or "null" context. That is, sometimes the context provides us no reason to depart from the literal meaning of a sentence. In such cases, the interpretation of a sentence is exactly its literal meaning.

(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 having to do with 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 implicature, conventional

implicature, the use of indirect speech acts, and the like.

- (a) **The Primacy of Core Grammar** - Since sentence meaning is purely compositional, the only linguistic conventions pertinent to it are those of "core grammar" and the lexicon. To say that the speaker's meaning of an utterance 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) **Literal Meaning is Sometimes Speaker Meaning** - 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 Null Context assumption, in which (i) one interpretation of the sentence is its literal meaning, and (ii) the speaker wishes to communicate exactly this content.

- (c) **The Primacy of Literal Meaning** - 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 orthodoxy, all the various interpretations we may wish to associate with a sentence other than its literal meaning are not part of 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.

2.1. The Left-Wing Challenge

This right wing orthodoxy has been subject to radical challenges from the intellectual "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 (1978). Searle claims that either we must abandon either the idea that literal meanings establish truth conditions, or the notion that literal meanings can be determined independent of context. To show that these two tenets of the orthodox position are in conflict, Searle considers the interpretation of simple sentences for which the existence of a literal meaning has generally gone unquestioned, such as the following:

(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 purported 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, one can respond to the owner's question about the location of

the cat by saying (1). Searle claims that in this situation, this utterance “should probably be described as an ingenious lie”. However, the use of this same sentence to describe the same situation, but in a different context, might have quite a different truth value. Thus, if both the speaker and hearer are aware of the unusual orientation of the mat, and the hearer is trying to determine on which of a number of possible objects the cat might be, then uttering (1) would be perfectly valid.

The same sentence seems to make different contributions to truth conditions in the two contexts, even though it describes precisely the same cat-mat relation. Searle concludes that either we must give up the idea that there is a completely context-independent literal meaning for all sentences, or give up the idea that such a literal meaning determines the truth conditions of an utterance.

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 one ran the cake over with a lawn mower, say, in response to the speaker’s uttering (2), that hearer would not have interpreted the speaker correctly. Since taking a lawn mower to the cake would not comply with (2), and since “cut” is not being used ambiguously, then it must be the case that the literal meaning of these sentence must take context into account.

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

2.2. The Right-Wing Defense

Katz (1981) responds to 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 dogma 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 radically 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. Indeed, someone who feigned compliance with (2) by snipping a blade of grass with a nail clipper might be accused of interpreting the

speaker "too literally". Whatever literal meanings might be, they are not meant to determine the truth conditions of the use of sentence as an utterance. Instead, the actual contribution of sentence meaning to utterance meaning is left unspecified by the orthodox position. Thus the sort of objections Searle raises cannot harm that doctrine.

2.2.1. A Note on Polysemy

A possible complication to these arguments is that words like "cut" and "on" are highly polysemous. If each different use of these terms in the sentences above corresponded to a distinct sense, then much of this argument would lose its force. Sentences like (1)-(3) would merely be highly ambiguous. To make the argument pertinent, we would need to show that the full spectrum of possible interpretations is not predicted by the lexicon alone.

Fortunately, it seems that, however specifically we are willing to postulate word senses, interpretations will be more specific still. For example, while we might postulate a specific sense of "cut" that means "slice", our interpretation of "Cut the salami" and "Cut the cake" will be rather different, even though the same sense seems to be in play. Thus, acknowledging polysemy affects where precisely the argument should be exploited, but does not endanger its essential force.

3. The New Problem

While Katz may effectively counter the details of Searle's argument, he 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 upon" and "balancing upon", 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” is neutral with respect to a great number of physical relationships, it appears as if “on” can be used to mean only one of these more specific relations. 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 upon the mat, attached to the mat in the manner of a notice on a bulletin board, or suspended from an inverted mat in the manner of a light fixture suspended from the ceiling.

To begin with, this simple fact undermines an important tenet of the right wing orthodoxy. This is the Null Context assumption, which states that the sentence meaning is utterance meaning in the null context. According to our argument, sentences like “The cat is on the mat” have a perfectly well defined sentence meaning, computable from the grammar and lexicon. This meaning is akin to that of “The cat is somehow supported by the mat.” But even in the null context, the meaning of this utterance is something akin to “The cat is lying upon the mat”, which is something else again. Thus, Null Context has been violated.

In fact, this argument suggests 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 does not confer meaningfulness upon it. The mistake is to assume that, because we heard a sentence in isolation (i. e., the null context), we are computing a meaning based only on linguistic knowledge. Even though 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. Thus, knowledge of the world suggests that “Cut the cake” refers to slicing, even if no further context is supplied. In contrast, the object computed from merely the linguistic facts, involving some less differentiated notion of cutting, may in and of itself not even be communicable.

Most importantly, the assumption of the Equivalence of Sentence and Literal Meaning has been violated. If the sentence “The cat is on the mat” is being used to communicate that the cat lying on the mat, it certainly does so *literally*. However, this “literal meaning” is not the same as the privileged “sentence meaning”, which is much more abstract.

Indeed, the term "literal meaning" is entirely inappropriate throughout all these discussions of sentence meaning. 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 meaning of a sentence. Thus, it seems that the term "literal" is useful only for talking about interpretations, not about intrinsic properties of sentences.

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, a sentence may mean something literally that is 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 (that is, something that can be meant); "literal meaning" is something radically different from sentence meaning (it is something that can be meant, but which may depart from the sentence meaning in important ways).

4. The New Picture

These challenges to the orthodox position seem to dictate the following picture: There is an object we can assign to a sentence, independent of context. This object pertains to the meaning assignable to usages of this sentence, but may not itself be a meaning per se. This object, together with what a language user knows about the world and the current context, contributes to an interpretation of the sentence. This interpretation is separate from speaker meaning, and logically prior to any form of rhetorical device, e. g., indirection, irony, etc.

Let us introduce a terminology to help describe this situation. We will call the semantic object assignable to a sentence in isolation the sentence's *primal content*. This object generally will be interpreted in a variety of fashions to become a meaning, a mapping involving world knowledge and situational context. This meaning, which a use of the sentence might convey, is called the *actual content* of the utterance. Note that actual content is a notion applicable to sentence use. However, we can use the term to talk about a

sentence, meaning the actual content one would associate with that sentence, given no additional contextual information, but assuming some body of world knowledge. Thus, when we say that the actual content of the sentence "The cat is on the mat" refers to lying down, we mean that this is the likely actual content to be assigned a use of this sentence given only our knowledge of the world and no additional knowledge of the situational context.

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.

4.1. Idiomaticity and Indirect Speech Acts

Above I have argued that it might be sensible to distinguish what we can assign to a sentence using only grammatical and lexical knowledge, from what we may assign to a sentence (or its use) by supplementing this information with world knowledge. Since "literal meaning" seems to involve interpretations of the latter sort, it cannot be identified with former.

From this perspective, let us re-examine the orthodox tenet of "Primacy of Core Grammar". Recall that this is the right-wing assumption that brutally distinguishes "core" grammatical knowledge from less generative linguistic conventions. Now, the entire motivation for this assumption in the first place appears to be that idiomatic conventions and the like are by definition not contributive to "literal meaning"; since literal meaning and sentence meaning are supposed to be the same thing, such conventions cannot pertain to sentence meaning.

However, once literal meaning is divorced from sentence meaning, there appears to be no motivation whatsoever for such a separation. Rather, it seems unavoidable that we include all linguistic constructions in the domain of grammatical knowledge when talking about computing something from a grammar and lexi-

con. Thus, in addition to the quite general facts about common syntactic units, one would also include as grammatical knowledge more specialized constructions (of the sort needed to comprehend "The more, the merrier") as well as (possibly syntactically redundant) facts needed to recognize certain forms as having idiomatic interpretations. (Systems such as that described in Wilensky and Arens (1980), and the *construction grammar* approach described in Fillmore, Kay and O'Connor (1987), Fillmore (1987) and Lakoff (1984) are developed along such lines.)

According to this account, sentences such as

(6) John kicked the bucket.

are *grammatically* ambiguous (rather than having grammatical sentence meaning and an extra-grammatical 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.

As another example, consider the treatment of sentences such as the following:

(7) Why don't you sit down?

In a formulation that identifies sentence and literal meaning, we would have little choice but to attribute to this sentence the literal meaning of being a question about why someone does not sit down, and, indirectly, a invitation to do so. Such cases are troubling because the subsequent interpretation cannot be made based on the literal meaning alone, and because computing the literal meaning is of dubious psychologically status. However, it is possible to postulate a construction that subcategorizes why-questions, and associates with such a subcategory the interpretation as an invitation. Indeed, the existence of this construction is supported by the observation that the form "Why don't you *be* seated?" is not a syntactically acceptable question.*

Including idiomatic interpretations and some indirect speech act forms as grammatical choices seems to have some empirical support. Gibbs (1983, 1986) offers evidence that strongly conventionalized linguistic

*I am indebted to Charles Fillmore for this observation.

forms are interpreted non-literally **without any alternative literal interpretation** being computed. Thus hearers are apparently prone to perceiving “kick the bucket” as meaning “die”, without typically computing a content pertaining to kicking. A similar phenomenon seems to occur with highly conventionalized speech acts, such as “Can you pass the salt?”. Indeed, hearers seem strongly biased towards computing the “conventional meaning” of an utterance, even when the context biases one toward a literal interpretation (Gibbs 1982).

This being the case, it seems most reasonable to treat the choice between a conventionalized, non-literal interpretation and a core-grammatical, literal interpretation as essentially the same as a choice between two core grammatical or two lexical interpretations. In particular, the choice of one alternative seems to inhibit consideration of the other.

I mention in passing that the inclusion of conventionalized expressions on par with core grammatical knowledge does not imply an impoverished representation of the former. For example, the fact that there may be specialized knowledge that lets one associate “Can you pass the salt?” directly with its indirect interpretation says nothing about whether additional facts about this construction are encoded. For example, one’s knowledge may contain a representation for the fact that this construction has its meaning by virtue of a quite general principle that would enable a hearer not familiar with this expression to nevertheless interpret it correctly, although perhaps by performing a greater amount of computation. Such knowledge may not be accessed in ordinary understanding situations, but may have empirical consequences in others.

4.2. Conventionalized Metaphor

Not all conventionalized forms are meant to be included directly as constructions. In particular, analyses that make recourse to conventionalized metaphor (e. g., Lakoff and Johnson 1980, Jacobs 1985) do not, but fit in nicely within this framework nevertheless. Consider a sentence like

(8) John gave Mary a kiss.

If we hypothesize that a conventionalized metaphor underlying this sentence were synchronically present,

then it is possible to treat this sentence as having a primal content that involves transferring, but an actual content that involves kissing. In such cases, it is a conventionalized metaphor that supplements the primal content to produce the actual.

Similarly, we may postulate a conventionalized metaphor underlying sentences like the following:

(9) John and Mary buried the hatchet.

If so, then the primal content of (9) refers to hatchet burying, and the actual content, to making peace.

4.3. Inferring Actual Content

A process that understands sentences in context is faced with the task of computing an actual content given a relatively underspecified linguistic input and an in-principle unbound quantity of world knowledge. However, the analysis offered here does *not* specify that a primal content be computed prior to or independently from an actual content. Rather, these terms refer only to the components of the semantic representation of an utterance that one might distinguish based on the kinds of knowledge necessary for computing them.

While actual/primal content is not a processing distinction, it still may be useful to distinguish those ways in which actual content may differ from the primal content of an utterance. It seems as if there are two kinds of differences: An actual content may be a more precise notion than the primal content, or, the actual content may involve a reinterpretation altogether, using some conventionalized mappings (such as a conventionalized metaphor).

I have previously referred to the first distinction, that of computing a more precise interpretation from a vaguer semantic content, as *concretion* (Wilensky 1983). In addition, in Wilensky (1986), we introduced a device called *views* for representing structure mappings like conventionalized metaphorical knowledge. For example, one can talk about viewing acting upon an object as transferring that action to that object. Such a view might be said to underlie sentences like (8) above. Since using views to help interpret an

utterance involves finding a structure that the given object may be viewed as, we call this distinction *unviewing*.

I emphasize that both unviewing and concretion do not necessarily represent independent inferences processes. Indeed, I think such independent processes to be unlikely. Norvig (1986) presents an inference process that makes these inferences, and several other types as well, but without resorting to separate processes for each type of inference. It would appear to be a straightforward extension of this work that the same mechanism that used knowledge to disambiguate a word (a process necessary for determining primal content) also performed concretion (necessary for actual content).

5. Other Objections to Literal Meaning

A number of other objections that have been raised against the notion of literal meaning. These are objections to the assumption that the literal meaning of a sentence plays a role in determining the speaker meaning of an utterance. It is useful to examine such arguments in the light of the preceding discussion. In particular, we need to ask whether such arguments, to the degree they are valid, might also be valid objections to the primal/actual content distinction offered here.

For example, researchers have objected to the notion that literal meanings are in some sense privileged by pointing out some common idiomatic expressions (most notably, "by and large" and "trip the light fantastic") have no acceptable literal meanings. Since such expressions have perfectly acceptable meanings, it must be the case that literal meaning is not always computed to attain speaker meaning.

First, note that such an argument is ineffectual against the primal/actual content distinction. This is because the position advocated here differentiates the literal-idiomatic distinction from other literal/non-literal contrasts. In particular, idiomatic constructions contribute to primal content on the same basis as do non-idiomatic constructions. That is, the formulation offered here assigns a primal content directly to constructions like "by and large". Such an assignment need not make any reference to the meanings of the constituents of such an expression. So such expressions always have a primal content, albeit a non-literal

one.*

In any case, it is important not to overextend this argument. Since literal meaning has been equated with sentence meaning in the orthodoxy, acknowledging that some sentences do not have literal meanings in the literal/idiomatic sense has generally been interpreted as acknowledging that they have no sentence meaning (in the sentence meaning/speaker meaning sense). However, the primal/actual content distinction avoids having to make such a conclusion. Also, the argument cannot be extended directly to other forms of non-literality. In particular, it says nothing about the utility of this distinction for direct/indirect speech acts and the like.

However, other arguments to this effect have been made. For example, consider the argument that Levinson (1981) advances to contradict the claim that the literal force of an utterance must play a role in determining the nature of the speech act considered. His argument is that sentences such as "May I remind you that your account is overdue?" cannot have a literal force because the reminding act is performed without the requested permission being granted.

Another example offered as evidence against the existence of a literal meaning when a non-literal interpretation is intended is the sentence "Would you mind telling me what time the store closes?" As Clark (1979) has pointed out, merchants responding to this question often respond by saying "Yes, we close at six." This has been taken as evidence that the literal meaning of the utterance once again did not determine the nature of the speech act involved.

Now, it would not be problematic to accept these objections and account for them within the primal/actual content framework. To the extent that this argument is correct, we could identify such utterances as conforming to special constructions, and attach their indirect sense as their primal content. This is precisely the treatment advocated above for utterances like "Why don't you be seated?" However, the primal/actual content framework does not require such treatment. The issue is in principle an empirical one, and the traditional distinctions are not always incorrect. I believe such is the case for these examples.

*In such cases, the actual content may or may not be identical with the primal content.

In the case of the first sentence, it seems possible to preserve a sentence/speaker meaning distinction simply by construing the speaker's intentions as insincere. More specifically, one might grant the sentence itself a meaning, namely, a request for permission to do an action, and interpret the speaker's action as a polite but insincere request for permission which happens to have the pragmatic consequence of effecting that for which permission is requested. In this case, the sentence meaning/speaker meaning distinction would coincide with the literal/non-literal distinction. Such analysis seems to be necessary in any case in order to account for the unusual nature of this sentence.

Thus, the example does not refute the claim that the sentence has a meaning that contrasts with the meaning of its use. It merely shows that the use of the sentence violates a felicity condition. Indeed, it is possible to respond to the sentence by saying "Yes, thank you" (in the sense of the "yes" granting permission rather than acknowledging the action), thus paying homage to the literal content.

Similarly, the fact that many (although by no means all) responders to "Would you mind ..." questions begin with "yes" seems to say little about the role that the meaning of such sentences might play in fashioning such a response. The fact that "yes" is not a response to the sentence meaning does not mean that the sentence meaning plays no role in shaping this speak act. One analysis that preserves the sentence/speaker meaning distinction is to assume that both the direct and indirect interpretations are available when formulating a response. Thus, the hearer represents that "Would you mind..." is a question being used as a method for requesting "Would you please..." The latter is further encoded as a question that serves to make a request. In this analysis, responses that begin with "yes" are making recourse only to the indirect part. That is, it is as if the speaker were asked "Would you tell me..." However, in response of the form "No,..." or "No, not at all....", the "no" refers to the direct question and the rest to the indirect material. Again, the fact that some speakers do indeed respond with "No,..." or "No, not at all...." appears to acknowledge the direct content of the utterance sometimes plays a role. The positive or negative choice would then appear to be an individual difference, depending upon whether the responder wished to address the direct question, or the indirect request.

5.1. An Example

The point of the preceding discussion is to demonstrate that there is considerable room to maneuver within the primal/actual content framework. In particular, I assume that a quite complex construal of a give sentence forms the basis for an actual response. This construal typically will contain elements corresponding to primal content, actual content, indirect speech acts, and a rich goal analysis.

As an example that involves both primal and actual contents and speech act interpretation, consider the following:

(10) Would you mind opening the window?

The primal content of this sentence includes that it is a question, and that the content of the question asks the hearer about the hearer's willingness to open a window. The level at which "open the window" is represented is rather abstract. According to one analysis, the verb "open" contributes to representation something akin to "cause an opening to be created by acting upon an object" (in this case, a window). This representation is neutral with respect to whether the opening is created in the object, as in "open the package", or in some other object in which the object acted upon is some sort of portal, as is the likely intention here, or in either, as in "open the drawer". Thus, from the primal content alone, one could not determine whether one were begin asked create an open in the window, or create an opening is something else, say an associated room or wall.

The actual content would represent this distinction, however. Here "opening the window" would be further concreted to something roughly equivalent to "create an opening by moving the window in manner it was so designed". This whole actual content (including the fact that it is a question) then might be interpreted as indirectly making a request, as was suggested in the previous example. This request to open the window (in the normal fashion) would probably be further analyzed at the level of goal analysis to be interpreted as a plan for the speaker's goal of getting a window opened for some contextually determinable purpose and length of time. For example, a hearer who, in response to this utterance opened the window in question and then immediately closed it again might be said to have fulfilled even the indirect request to

open a window, but to have ignored one of the likely underlying speaker goals.

It is worth noting that there are several analyzes other than this one that are compatible with the overall paradigm. For example, it could be the case that the lexicon distinguishes "open <portal>" from "open <container>", so that the primal content of "open the window" was something akin to "create an opening by using the window", and that of "open the jar" akin to "create an opening in the jar." In this analysis, the actual content is still more specific. For example, the actual content of "open the jar" is likely to include the equivalent of "by unfastening its lid", and "using the window" will get concreted to include "in the conventional fashion", so that these requests could not be complied with by firing at the objects in question with a canon, etc.

5.2. Psychological Considerations

Another set of objections challenges the orthodoxy as a psychology model. For example, consider experiments of the sort described in Gibbs (1984). The basic structure of such experiments is to situate in context both utterances that should be interpreted literally, and those that should be interpreted in some other way (say, as indirect speech acts). Then a reaction time is taken for the subjects to make a paraphrase judgment on the target sentence. In general, in context, the non-literal interpretations take no longer to process than literal ones, although subjects do indeed take longer to read and make paraphrase judgments of "non-literal" sentences out of context. Such results are generally taken as evidence that the literal meaning is not necessarily computed in context, since doing so would require an additional step when a non-literal interpretation is subsequently sought.

As mentioned above, the primal/actual content distinction addresses this criticism when "literal" is contrasted with "idiomatic". Moreover, in some cases, I propose encoding what might otherwise be construed as an indirect speech act as the primal content associated with a specific construction. However, in those cases in which one might want to postulate a content for the sentence related to the content of the utterance through some manner of indirection, the same sort of criticism may be applied to the framework suggested here. That is, the actual content of a sentence might be used to invoke a rule that suggests an

indirect speech act interpretation of that sentence used as an utterance. Similarly, the treatment of metaphors postulates the existence of a "literal" rendering and an actual, metaphorically interpreted one. In such cases, we can ask whether such renderings are actually computed by a hearer.

While I am explicitly not advocating a processing theory that specifies the order of computation of primal content and actual content, I believe that the psychological results have been interpreted overly negatively with respect to whether multiple levels of interpretation can play a role in sentence use. The primary reason is that accepting a particular interpretation of a sentence without needing to compute an additional interpretation may involve just as much work as computing the additional interpretation.

To see why this is so, first consider those cases in which a high degree of conventionality may be involved. Conventionalized metaphors and some indirect speech acts are cases in point. In these situations, a system must check to see whether there is a conventionalized metaphor or an applicable speech act rule, and then apply the same. The catch is, one must perform this checking in all circumstances. Moreover, one may find and apply rules to a sentence, and then reject the resulting interpretations. In other words, the amount of time spent in deciding to accept the interpretation of some utterance may be no different than the amount of time needed to compute and accept an additional interpretation of that utterance.

A similar argument holds for less conventionalized cases. In these cases, we may first judge an utterance defective in some manner before attempting to compute additional interpretations. But as before, the judgment that an utterance is defective must be done even if the initial interpretation of the sentence is to be accepted. For example, an utterance intended metaphorically might first be judged defective before an attempt is made to understand it as a metaphor. Similarly, the utterance intended as an indirect speech act might first have to fail to make sense as a direct speech act. But to do so, the process of determining whether the utterance is defective must be done in every case.

Now, such processes are not at all well understood. But such processes apparently require recourse to all sorts of world knowledge, and require the hearer to make judgments about "appropriateness" and other rather open-ended concepts. Very likely, the time required for such judgments greatly exceeds that for

computing the initial interpretation of the utterance, a far more constrained task. Indeed, such processing may be performed so that computing such interpretations "falls out" of the general process of determining the appropriateness of the initial reading, so that little or no additional time is required to compute such an interpretation if the literal meaning is found defective in some way. Additional interpretations, e. g., metaphoric ones, would be computed gratuitously.

Furthermore, since the process of judging appropriateness of an interpretation is essentially one of taking contextual considerations into account, contextual biasing effects would strongly favor and facilitate computing one interpretation rather than another. Thus, a listener having witnessed an obviously stupid action might expect a sarcastic comment; in the same situation, a literal comment may be far less apt. Thus it should not surprise us if such expectations facilitate computing non-literal interpretations to the point where they would be more quickly accepted than a less motivated literal one.

Of course, these comments are not meant to be interpreted as an argument for the psychological validity of the computation of various levels of meaning. They merely suggest that the current evidence about its validity is scant indeed when one considers the overall model of language processing in which such a theory must be embedded.

5.3. Vagueness

Searle (1980) makes several objections against a position similar to the one I advocate here. In particular, he considers the possibility that sentences like (2) and (3) are simply vague, and that "the contexts of the rest of the sentences ... enable the hearer to infer what the speaker means even though the speaker meaning was not precisely expressed by the literal meaning of the sentence uttered...." He then raises three problems for this position: First, he claims that this position predicts we should understand the meaning of sentences like

(11) Sam cut the coffee.

while in fact, we cannot. Second, "crazy misunderstandings" would count as literally correct

interpretations of ordinary sentences. Finally, Searle argues that it seems most implausible that we go from a common literal meaning of sentences like (2) to different speaker meanings.

Let us examine each objection in turn. First, Searle claims that the vagueness position should predict that we understand sentences like (11), since we understand each of the words. But this is simply false. The vagueness position, at least in the primal/actual content interpretation, doesn't guarantee that, if the vague (i. e., primal content) interpretation gives us truth conditions, that we will know when those conditions are met. Thus, if the primal content of (11) is something like "Sam made a division in the coffee", we are by no means guaranteed that we will know how coffee could be divided. Our *understanding* of what the speaker meant would require this, but the semantic content of the sentence does not. Certainly, the sentence would be literally true if we somehow divided the coffee. It's just not clear how to do this.

Moreover, the situation doesn't improve if the vagueness goes away. Consider the following sentence:

(12) John cut the coffee in exactly the same manner that he cut the cake.

Presumably, cutting the coffee here is just as puzzling as it was in (11). But here, the interpretation of "cut" appears to be just as precise as it is in (2) or (3). So the relative specificity appears not to be a contributing factor.

Furthermore, consider a sentence like the following:

(13) John foolishly tried to cut the coffee, but of course he failed.

This sentence seems perfectly reasonable. The truth conditions for cutting weren't met, but it is precisely because the first clause has such truth conditions that we can understand this sentence. Thus, inability to determine how to meet truth conditions does not impugn the existence of such conditions.

Part of the problem here is that Searle assumes "literal meaning" really must be a meaning, and it is hard to see how a real meaning (i. e., something we may have intended to communicate) could exist for a sentence which we are inclined to say we don't understand. But the primal/actual content distinction resolves

this difficulty: The primal content need not be an actual meaning, while our understanding must be. It may be difficult to determine an actual content from a primal one. But that in itself does not refute the latter's existence.

Searle's second argument is that "crazy misunderstandings" would have to be construed as literally correct interpretations of a sentence. He offers as an example the case in which a contract is made for someone to cut someone else's grass on a weekly basis, but the individual so contracted attempts to comply with the contract by all sorts of unorthodox cutting actions. Searle claims that such an individual would not have complied with the letter of the contract. He made be right, but, as Katz (1981) points out, compliance with a contract is a language use issue. Certainly, our various conventions of language use (not to mention contractual agreements) would dictate a particular interpretation of cutting the grass. But this says nothing about whether a vague notion of cutting, i. e., a primal content, is assignable to the sentence, and perhaps even contributed to this understanding. Recall that the primal content is not supposed to determine the conditions of satisfaction of an utterance, but merely the truth conditions of a sentence. "Crazy misunderstandings", as suggested above, often merit the response that the speaker was taken too literally.

Finally, we have the argument that computing the specific meaning from the vague one via some inferential process seems "implausible". This objection is simply irrelevant, because, as stated above, the primal/actual content distinction is not a processing distinction, but a distinction based on what one can say about a sentence given certain kinds of knowledge. In fact, I suspect that pragmatic knowledge, contextual priming, etc., enters into understanding processes at the earliest stages, and that one does not proceed from vague to specific interpretations, at least in all cases.

In sum, to the extent that the primal/actual content position may be construed as a vagueness position, the position seems well warranted.

6. Summary

There is much to be said for distinguishing the various kinds of knowledge that are necessary to achieve

mastery of a language. In particular, it seems reasonable to ask whether some fact is a fact about a conceptual system in general or a fact about language per se, whether a convention is a central core construction of a language or a peripheral, derivative one, whether an analysis of a sentence makes recourse to only linguistic facts or to other facts as well, and whether there is psychological relevance to any of these distinctions.

However, when it comes to grouping such knowledge, this has often been done erroneously. In particular, it has been assumed that the interpretation of a sentence out of context would generally not involve extralinguistic knowledge, whereas in fact, knowledge about the world is typically brought into play by default. It has been assumed that, since the term "literal" distinguishes productive uses of words from idiomatic uses, non-metaphoric from metaphoric, and direct from indirect, that literal meanings must be the same as sentence meanings, i. e., that they could be computing from knowledge of the words and core grammar rules of the language, and moreover, that such a distinction is a pertinent one. In fact, while the term "literal" does make these distinctions, it does not necessarily separate interpretations which involve context and world knowledge from those that do not. In particular, the productive versus idiomatic aspect of the literal/non-literal distinction operates within the domain of linguistic convention per se, while both the literal and non-literal interpretations of a sentence typically are made involving recourse to knowledge of the world.

Since the terms "literal meaning", "sentence meaning" and "speaker meaning" are used in a manner that promotes erroneous groupings, a new set of terms is offered. "Primal content" refers to the interpretation we can assign to a sentence based on lexical and grammatical knowledge, broadly construed. In particular, both productive and non-productive constructions are included under this banner. In addition, the resulting interpretation may in and of itself not be a meaning, in that it may be an unsuitable candidate for communication. In contrast is "actual content", which refers to the specific meanings speakers encode into utterances and extract out of utterances, generally making liberal use of everything they know. Actual content is by no means the end of the meaning story, as actual contents may encode only a portion of a speaker's intention, or in fact belie it altogether.

While this distinction is offered as a more useful framework to talk about the meaning of sentences and utterances, it remains only a theoretical construction. In particular, no claim is made about the utility of an autonomous computation of primal content, or of the psychological reality of doing so. The distinction is meant to serve as a framework for organizing the knowledge that participates in such tasks, and for describing the contributions of parts of this knowledge to the whole.

7. Acknowledgments

I would like to thank Professor Raymond W. Gibbs, Jr. for his comments on an earlier draft of this paper.

8. References

- (1) Clark, H. Responding to indirect speech acts. *Cognitive Psychology*, 11, pp. 430-277, 1979.
- (2) Fillmore, C., *On Grammatical Constructions*. Forthcoming, 1987.
- (3) Fillmore, C., Kay, P., O'Connor. Regularity and idiomaticity in Grammatical Constructions: the case of *let alone*. University of California Cognitive Science Program Report No. 48, June 1987.
- (4) Gibbs, R. A critical examination of the contribution of literal meaning to understanding nonliteral discourse. In *Text 2* (1-3), pp.9-27. 1982.
- (5) Gibbs, R. Do People Always Process the Literal Meanings of Indirect Requests? In *Journal of Experimental Psychology: Learning, Memory, and Cognition*, Vol. 9, No. 3, pp. 524-533. 1983.
- (6) Gibbs, R. Literal Meaning and Psychological Theory. In *Cognitive Science*, vol. 8, no. 3, pp. 275-304. 1984.
- (7) Gibbs, R. Skating on thin ice: Literal Meaning and Understanding Idioms in Conversation. In *Discourse Processes*, 9, pp.17-30. 1986.
- (8) Jacobs, P. S. A knowledge-based approach to language production. Report no. UCB/CSD 86/254. Computer Science Division (EECS), UC Berkeley. August, 1985.
- (9) Katz, J. Literal meaning and logical theory. *Journal of Philosophy*, 78 pp. 203-234. 1981.
- (10) Lakoff, G. There-constructions: A case study in grammatical construction theory and prototype theory. University of California Cognitive Science Program Report No. 18, January 1984.

- (11) Lakoff, G. and Johnson, M. **Metaphors We Live By**. University of Chicago Press, Chicago, Ill. 1980.
- (12) Norvig, Peter. **A Unified Theory of Inference for Text Understanding**. UC Berkeley Computer Science Report No. UCB/CSD 87/339. January, 1987.
- (13) Searle, J. **Literal meaning**. In J. Searle (ed.). **Expression and Meaning**. Cambridge, England, Cambridge University Press. 1979.
- (14) Searle, J. **The background of meaning**. In J. Searle, F. Keifer, & M. Bierwisch (eds.), **Speech Act Theory and Pragmatics**. Boston, Mass., D. Reidel. 1980.
- (15) Wilensky, Robert. **Planning and Understanding: A Computational Approach to Human Reasoning**. Addison-Wesley, Reading, Mass. 1983.
- (16) Wilensky, Robert. **Some Problems and Proposals for Knowledge Representation**. Berkeley Computer Science Technical Report UCB/CSD 86/294, May 1986.
- (17) Wilensky, Robert and Arens, Yigal. **PHRAN - A Knowledge-Based Approach to Natural Language Analysis**. Berkeley Electronic Research Laboratory Memorandum No. UCB/ERL/M80/34. August, 1980.

