The alleged priority of literal interpretation

François Recanati

The alleged priority
of literal interpretation

François Récanati

CREA, CNRS/Ecole Polytechnique, Paris

The alleged priority
of literal interpretation

In this paper I argue against a widely accepted model of utterance interpretation, namely the LS model, according to which the literal interpretation of an utterance (the proposition literally expressed by that utterance) must be computed before non-literal interpretations can be entertained. Alleged arguments in favor of this model are shown to be fallacious, counterexamples are provided, and alternative models are sketched.

The alleged priority
of literal interpretation*

1. Introduction: the standard model

According to the standard model (Grice 1989), the interpretation of non-literal utterances proceeds in two stages: [a] the hearer computes the proposition literally expressed by the utterance; [b] on the basis of this proposition and general conversational principles, he or she infers what the speaker really means. In what follows, I will try to show that this model, which is still accepted by a majority of researchers in semantics and pragmatics, is

unfounded and deserves questioning. More specifically, I will argue that the so-called 'proposition literally expressed by an utterance' - assuming that there is such a thing - need not be computed in the process of interpreting a non-literal utterance.

Doubts have often been expressed in the cognitive science community as to the validity of the standard model (see e.g. Rumelhart 1979, Gibbs 1984, Sperber & Wilson 1986, Langacker 1987). In particular, a number of psycholinguistic studies have demonstrated the equivalence of processing times for the interpretation of literal and non-literal utterances in normal conversational settings (Harris 1976, Ortony et al. 1978, Gerrig and Healy 1983, Gildea and Glucksberg 1983, Inhoff et al. 1984, Gibbs 1986a). This casts doubt on the standard model, according to which the literal meaning of a non-literal utterance is processed first; for non-literal utterances should take more time to process than literal ones, if the model was correct. Despite these results, the standard model is still with us - it still is the dominant model. Among the reasons which may explain why this is so, two seem to me especially important and worth thinking about. First, as Gibbs and Gerrig 1989 pointed out, the sheer equivalence of processing times is not sufficient to actually *refute* the standard model - one might, for example, argue that a 'masking' process fosters the illusion of equivalence (Gibbs and Gerrig 1989: 150-1, 155). Here as elsewhere, processing time experiments, suggestive though they are, are not totally conclusive. On the other hand and more importantly, it is felt that the standard model *has got to be correct in some sense*. It seems to many that there is an a priori, philosophical argument in favor of the standard model.

In this paper, I want to attack the standard model from a different angle, by considering its philosophical underpinning, i.e. the sort of a priori argument I have just alluded to. Though the argument in question is seldom made explicit, I think it can be spelled out along the following lines:

Given what we might call the asymmetric dependence of non-literal meaning on literal meaning, the meaning contextually conveyed by an expression counts as 'non-literal' in the full sense of the term only if it is derived from some basic, literal meaning which must therefore be processed for the former to be accessed. Cases in which the alleged non-literal meaning can be accessed directly (without processing the literal meaning first) are cases where the non-literal meaning has got conventionalized in some way and is no longer 'non-literal' in the full sense. It is no coincidence, from this point of view, that so many of the psycholinguistic experiments purporting to refute the standard model (e.g. Clark 1979, Ortony et al. 1978, Gibbs 1983, 1986b) were actually concerned with idioms or conventionalized indirect speech acts such as 'Can you pass the salt' .

To many that sort of argument seems powerful and convincing. This goes a long way towards explaining why the standard model is still considered adequate, psycholinguistic evidence to the contrary notwithstanding (see e.g. Dascal 1987). But I will show that the 'asymmetric dependence' type of argument is fallacious and cannot be invoked in defence of the standard model. (For those who are already convinced that the standard model is mistaken, let us say that I want to nail its coffin.) I will proceed by sketching alternatives to the standard model and showing that they are not threatened by the type of argument in question. At the end of the paper I will turn to more empirical issues and indicate a class of examples which the standard model simply *could not* handle - a class of example which, more directly perhaps than processing time experiments, shows the standard model to be false.

2. Varieties of contextual interpretation

Comprehension often involves contextually assigning a particular value to some constituent of the uttered sentence. I use 'contextual interpretation' to refer to the process whereby a semantic value is contextually assigned to an expression. Non-literal interpretation is a particular case of contextual interpretation, namely the case in which a constituent of the uttered sentence is contextually assigned an interpretation distinct from its literal interpretation, e.g. a metaphorical one or a metonymical one. There are other sorts of contextual interpretation, however, and I believe systematic comparison between them can shed considerable light on non-literal interpretation itself. In this section I present four types of contextual interpretation, distinct from the non-literal type. In later sections I will introduce various models of contextual interpretation and discuss their respective applicability to the various types.

• *Sense selection*. When the semantic value to be contextually assigned is one of the conventional 'readings' of the expression, the contextual process of value assignment is called 'disambiguation' or 'sense selection': An expression has more than one conventional sense, and comprehension involves selecting one of the senses as contextually appropriate. There are many cases in which the semantic value contextually assigned to a constituent of the sentence is not linguistically conventional, however. All the types of contextual interpretation to be mentioned below fall into the latter category.

• *Sense construction.* Often, the conventional sense of the constituents of a complex phrase and the way they are grammatically combined is not sufficient to determine the semantic value of the complex phrase involving those constituents; the overall meaning of the complex phrase must be constructed in context. Thus 'He finished the book' can mean that he finished reading the book, writing it, binding it, tearing it into pieces, burning it, and

so forth (Pustejovsky 1991, Cohen 1985, Langacker 1984); 'finger cup' will mean either *cup having the shape of a finger* or *cup containing a finger of whisky* or *cup which one holds with one finger*, or whatever (Bühler 1934, Kay and Zimmer 1967); 'John's book' can mean *the book that John owns*, *wrote*, *gave*, *received*, or whatever (Kempson 1977, Recanati 1989). In all such cases there is not a 'selection' from a limited range of preexisting interpretations for the complex phrase. Rather, an indefinite number of possible interpretations can be constructed in a creative manner (Clark 1992).

• *Specifization*. This is the process by which a term which has a general meaning can be contextually interpreted in a much more specific sense (Atlas 1989; Bach forthcoming). For example the mass term 'rabbit' will be preferentially interpreted as meaning *rabbit fur* in the context of 'He wears rabbit' and as meaning *rabbit meat* in the context of 'He eats rabbit' (Nunberg and Zaenen forthcoming).[1] Again it's not a matter of selecting a particular value in a finite set; with a little imagination, one can think of dozens of possible interpretations for 'rabbit' by manipulating the stipulated context of utterance; and there is no limit to the number of interpretations one can imagine in such a way.

One important difference between specifization and sense construction is that the latter, but not the former, can be analysed as involving a hidden variable (Partee 1984, Recanati 1989). For example, 'John's book' can be said to mean, *the book that bears relation* r *to John*, where the variable '*r′*' must be contextually instantiated for the complex phrase to possess a definite semantic value. This analysis stresses the analogy between sense construction (as opposed to specifization) and another process to be contrasted with sense selection, namely reference assignment.

---

[1] The term 'specifization' comes from Bach. Nunberg & Zaenen use 'precision' in the same sense.

• *Reference assignment.* The semantic value of a referential expression is not the linguistic meaning of that expression (its 'character', in Kaplan's terminology) but the reference which must be contextually assigned to the expression by virtue of the semantics of the language (Kaplan 1989, Recanati 1993). The linguistic meaning of the expression only helps to contextually identify the reference. It follows that the semantic value or 'content' of the expression (viz. its reference) is contextually variable and does not possess the constancy of linguistic meaning. 'He's in Paris' can mean that John, Bob or *virtually anybody* is in Paris. Here also, we have a variable which must be contextually instantiated for the utterance to possess a definite semantic value. A reference has to be found (and can be freely sought) in context for the pronoun 'he', in the same way in which a particular relation has to be found (and can be freely sought) in context in order to construct the sense of the complex expression 'John's book'.

There is much to be said about each of these processes (sense selection, sense construction, specifization, and reference assignment), and especially about the general picture of interpretation which emerges when they are taken seriously. In this paper, however, I will use them only to shed light on the fifth type of contextual process, namely non-literal interpretation. I will first consider two common models of sense selection (the serial model and the parallel model) and discuss their applicability to other types of contextual interpretation. A first version of the 'asymmetric dependence' argument, to the effect that only the serial model is applicable to non-literal interpretation, will be scrutinized and rejected. A parallel model of non-literal interpretation will accordingly be sketched. Next, I will draw a distinction between different types of serial model applicable to different types of contextual interpretation. The standard, Gricean model of discourse interpretation (the literality-based serial model, or 'LS model' for short) will be shown to be applicable to specifization and non-literal interpretation, but not to sense selection, reference assignment or sense construction. The only serial model applicable to the latter

processes of contextual interpretation is the accessibility-based serial model (AS model). I will argue that the AS model can also be applied to non-literal interpretation and specifization, in such a way that there are two alternatives to the standard, LS model. The gist of this paper is that both alternatives are actually preferable to the LS model.

3. Serial vs. parallel models

There has been a good deal of psycholinguistic research on sense selection, out of which two basic models have emerged:

• According to the **serial model**, the most accessible candidate for semantic value is tried first; if there is a problem some backtracking takes place.
• According to the **parallel model**, all candidates - or at least, all candidates which reach a certain level of accessibility - are tried in parallel; the first candidate whose processing yields satisfactory results in the broader context of discourse is retained, while the others are suppressed.

Before considering whether and how these two basic models apply to the other types of contextual interpretation, a few clarificatory comments are in order.

[1] By a 'candidate', I mean one of the semantic values which can be assigned to an expression in context. In the case of ambiguous words, the relevant semantic values - the relevant 'candidates' - are the (multiple) conventional readings of the expression. For example the candidates for semantic value in the case of an ambiguous word such as 'bank' are the two readings *river edge* and *financial institution.* But the notion of a candidate for semantic value extends beyond the case of ambiguous words; it applies whenever several distinct values can be contextually assigned to an expression. Thus the literal and the

metaphorical interpretation of an expression are two 'candidates'; different possible referents for a pronoun in a particular context are 'candidates', and so forth.

[2] Accessibility, for a mental representation, is an intuitive notion which can be made precise in many ways. According to Barsalou and Billman (1989), accessibility is a multi-factor phenomenon which involves frequency of processing, recency of processing and contextual relevance. Sperber and Wilson (1986) characterize accessibility in terms of recency of processing and/or associative links to recently processed information; they don't invoke relevance as they use accessibility in *defining* relevance. Following these authors I will assume that accessibility, for a mental representation, involves the following factors (including a 'recursive' one): recency of processing, close associative links to accessible representations, and frequency. As accessibility so characterized corresponds to the intuitive notion of 'salience', I will use both terms indifferently. Another expression I will use as equivalent to 'accessible' is 'activated'. Indeed, the three factors in terms of which I have characterized accessibility can easily be rephrased in the vocabulary of activation and activation spreading: Some representations are activated because they are being processed or have been recently (first factor); they spread activation to associatively related representations (second factor); and frequency of processing can be conceived of as lowering the activation treshold of the representation (third factor). My use of this vocabulary should not be construed as committing me to a particular framework, however. The notion of accessibility I have intuitively characterized can be embodied in many different frameworks, between which I do not feel compelled to choose, given the very general aim of this paper.

[3] The opposition between serial and parallel models should not be taken too strictly; it is, to a large extent, a matter of degree. If parallel processing comes with a bias, the parallel model does not significantly differ from the serial model (Garman 1990:360). Be that as it may I will not be concerned with the details here, but only with the general tendencies embodied in the two basic models.

The two models apply neatly in the case of reference assignment. On the parallel model, all (sufficiently accessible) candidates for the status of referent are tried in parallel in such a way that the first one which yields satisfactory results in the broader context of discourse is retained, the others being suppressed. On the serial model, the most accessible candidate is tried first and backtracking occurs only if there is a problem. Consider, for example, utterance (1) (from Recanati 1993:265):

(1) John was arrested by a policeman yesterday. He had just stolen a wallet.

In order to interpret the utterance one must assign a reference to the pronoun 'he' in the second sentence. The two persons who were mentioned in the first sentence, namely John and the policeman, are obvious candidates, and we may suppose that there is no other (sufficiently accessible) candidate. (There would be one if, for example, the speaker pointed to someone while uttering the pronoun.) According to the serial model, if one of the two candidates is more accessible than the other, it is processed first, and the resulting interpretation of the utterance is retained unless it turns out to be unsatisfactory. According to the parallel model, both candidates are processed in parallel, i.e. both interpretations (*the policeman had just stolen a wallet* and *John had just stolen a wallet)* are constructed and entertained until one is found satisfactory and the other is suppressed.

What about non-literal interpretations? Can we apply the two models in this domain also? In the literature only one model is considered seriously: the serial model. The parallel model is taken to be ruled out by the particular structure of non-literality, namely the fact that there is an *asymmetry* between the literal and the non-literal interpretation: the latter presupposes the former. This is so because the process of non-literal interpretation consists in *inferentially deriving* one interpretation from the other (Grice 1989). Now this means that non-literal interpretation proceeds *serially*. On the inferential model of non-literality,

we process the literal interpretation first, and go on to the non-literal interpretation when this is required to make sense of the speaker's utterance.

I agree that the inferential model of non-literality, which is accepted by so many philosophers and linguists, entails rejecting the parallel model of non-literal interpretation. For the inferential model of non-literality embodies a particular version of the serial model, which I call the LS model (the 'literality-based serial model', according to which the literal interpretation is processed first). Still, I think that the parallel model is *not* ruled out: there is still the possibility of *rejecting the inferential model of non-literality.* This is the line I will take in this paper.

4. A parallel model for non-literal discourse

What makes the serial model so attractive is the fact that the non-literal interpretation is derived from, and presupposes, the literal one. This I do not wish to deny - I admit that the literal interpretation comes first. Still, I want to resist the conclusion that the literal interpretation must be 'processed' first. To say that the literal interpretation is processed first is to say that we have a two-step procedure: (i) the interpreter accesses the literal interpretations of all constituents in the sentence and uses them to compute the proposition literally expressed; (ii) on the basis of this proposition and general conversational principles he or she infers what the speaker means (which may be distinct from what is said, i.e. from the proposition literally expressed). What I am disputing is not the claim that the literal interpretation of the constituent is accessed before the non-literal interpretation - that I take to be obvious - but the claim that the process of semantic composition which consists in putting together the semantic values of the parts to determine the semantic value of the whole begins by paying attention only to literal semantic values, and turns to non-literal values only after the literal semantic value of the whole (the proposition literally expressed) has been computed. It is this picture which I think is unwarranted.

If I am right, the asymmetric structure of non-literality (the fact that non-literal interpretations presuppose literal interpretations) does not rule out a parallel model such as the following:

Parallel model of non-literal interpretation
The literal meaning of the relevant constituent is activated, and this activation automatically (i.e. non-inferentially) spreads to various associatively related representations; the latter are potential candidates for the status of semantic value. All candidates, whether literal or non-literal, are processed in parallel and compete. When an interpretation which fits the broader context of discourse is found, the others are suppressed.

On this view, non-literal interpretations still proceed (associatively) *from* literal interpretations, which they indeed presuppose; but, although generated serially, they are processed in parallel. The literal interpretation has no privilege over non-literal interpretations; they compete and it is possible for some non-literal interpretation to be retained (if it fits the broader context of discourse) while the literal interpretation is suppressed. In other words, the non-literal interpretation is *associatively derived* from the literal interpretation, but it is not *inferentially* derived. Inferential derivation entails computation of the literal value of the global sentence (which serves as input to the inference), while associative derivation is a 'local' process which does not require prior computation of the proposition literally expressed (Recanati 1993:263-6).

Consider, as an example, Geoff Nunberg's famous ham sandwich (Nunberg 1979). The waiter says 'The ham sandwich has left without paying'. On the LS model the interpreter computes the proposition literally expressed by the sentence - namely the absurd proposition that the sandwich itself has left without paying - and from its absurdity infers that the speaker means something different from what she says. On the parallel model the

description 'the ham sandwich' first receives its literal interpretation, in such a way that a representation of a ham sandwich is activated; activation then spreads to related representations, including a representation of the man who ordered a ham sandwich.[2] *All these representations activated by the description 'the ham sandwich' are potential candidates for the status of semantic value of the expression;* all are equally susceptible of going into the interpretation of the global utterance. Now the ham sandwich orderer is a better candidate than the ham sandwich itself for the status of argument for '... has left without paying'. It is therefore the derived, non-literal candidate which is retained, while the literal interpretation is discarded.

An important difference between the LS model (according to which the literal interpretation is processed first) and the parallel model just outlined is this: on the parallel model is it possible for an utterance to receive a non-literal interpretation *without the literal interpretation of that utterance being ever computed*. The non-literal interpretation of the global sentence does not presuppose its literal interpretation, contrary to what happens at the constituent level. If the non-literal interpretation of some constituent fits the context especially well it may be retained (and the other interpretations suppressed) *before* the literal interpretation of the sentence has been computed. Whether or not this sort of thing actually happens, this is at least *conceivable*, on the parallel model. (On the LS model this situation is ruled out in principle.)

Of course I am not saying that the LS model never fits the facts - that we never compute the proposition literally expressed and infer what the speaker means. I fully admit

---

[2] Actually the issue is more complicated. According to Sag (1981) and to Nunberg himself in his most recent treatment of the topic (Nunberg 1993, forthcoming), it is not the description as a whole, but the predicate 'ham sandwich' *in* the description which has a derived, non-literal value; the ham sandwich example is a case of property transfer, not a case of deferred reference (see footnote 7).

that there are 'conversational implicatures' which work exactly that way (even though not everything which has been treated as conversational implicature in the literature on these subjects belongs to the same category [Sperber & Wilson 1986, Carston 1988, Recanati 1989, Bach forthcoming]). I am only saying that there is no conceptual necessity why we should (always) have to wait until the literal interpretation of the sentence is computed and contextually tested before trying non-literal interpretations. In particular, the fact that the non-literal interpretation of a constituent presupposes the literal interpretation of that constituent does not entail that the non-literal interpretation of the whole sentence also presupposes the literal interpretation of that sentence.

5. The accessibility-based serial model

So far I have mentioned two models for non-literality, the literality-based serial model (LS model) on the one hand and the parallel model. But another model deserves to be considered.

In the case of disambiguation and reference assignment also I made a distinction between two models, a serial one and a parallel one; but if one looks closely one sees that the two models distinguished for reference assignment and disambiguation do not coincide with the two models distinguished for non-literal discourse. The parallel model is in all cases the same, but the serial model envisaged for dealing with non-literality is quite different from the serial model envisaged for dealing with reference assignment and disambiguation. The former is a literality-based serial model, while the latter is an accessibility-based serial model.

Recall what was said concerning reference assignment. Whereas, on the parallel model of reference assignment, all (sufficiently accessible) candidates for the status of referent are processed in such a way that the first one to yield satisfactory results is retained, on the serial model the most *accessible* candidate is tried first. It is accessibility,

not literality, which matters for reference assignment, for one very good reason: the various candidates involved are on the same footing and cannot be discriminated in terms of literality. In example (1) above, John and the policeman can both serve as referent for the pronoun 'he' without either of them being more 'literal' than the other. It follows that, if one of the candidates is processed before the others, that cannot be because it is literal while the others aren't.

The same thing holds of ambiguity. In a process of sense selection the candidates are all literal interpretations of the expression - none is *more* literal than the others. It follows that if one is processed first, this can't be because it's the literal interpretation. In all those cases (reference assignment, sense selection, sense construction) we find the same pattern: the candidates are all on the same footing with respect to literality, so, if a serial model applies and one of the candidates is processed first, this cannot be because that candidate is literal (while the others are not). What determines that one candidate is processed before the other can only be accessibility - the fact that one candidate comes to mind more readily than the others.

To sum up, whenever the candidates are on the same footing with respect to literality, the only serial model available is the accessibility-based serial model (AS model), according to which the most accessible candidate is processed first. But when it is possible to distinguish one of the candidate as being 'literal', another serial model is available, namely the literality-based serial model (LS model), according to which the literal candidate is tried first. In such cases we have three models to consider: the parallel model, the LS model and the AS model.

Note that specifization and sense construction do not fall into the same category here. The candidates for the status of semantic value in a case of sense construction are on the same footing with respect to literality; whichever relation holds between John and the book in a contextual interpretation of 'John's book', that relation will not be more literal than another relation which would be relevant in a different context. The situation is very

similar to that of reference assignment. But in a case of specifization, there is one particular interpretation which is more literal than the others, namely the 'general' reading. The latter is more literal in the sense that it corresponds to what is linguistically encoded, while the contextually specifized readings go beyond what is encoded. Thus the mass term 'rabbit' literally means *rabbit stuff*. The other, more specific interpretations (*rabbit meat*, *rabbit fur*, and so forth), are in a sense less literal, whence it follows that a literality-based serial model ought to be available for specifization as it is for non-literal interpretation.

There is more to be said in this connection, but I cannot elaborate. The only thing that matters is that for non literal discourse as well as for specifization there are, in principle, three models to consider: the parallel model and *two* serial models, viz. the AS model and the LS model. All candidates may be processed in parallel, or the most accessible may be processed first, or the literal candidate may be processed first.

The parallel model and the AS model share a common feature as opposed to the LS model. On the LS model the literal interpretation of the sentence must be computed, whether or not it is the ultimately correct interpretation of the utterance. On both the parallel model and the AS model the non-literal interpretation of the global sentence does *not* presuppose its literal interpretation, contrary to what happens at the constituent level. The parallel model and the AS model have more in common than this negative feature, however; for the parallel model itself is based on accessibility. On the parallel model, the candidates which are actually processed are those that are sufficiently accessible; and the candidate which is ultimately retained as the correct interpretation is that which is most accessible when the interpretation stabilizes (§10). So interpretation is (almost) as much accessibility-driven on the parallel model as it is on the AS model. The real contrast, I think, is not between parallel and serial models, but between accessibility-based models and the LS model. Be that as it may, I am not arguing specifically in favor of either the parallel model or the AS model in this paper; I am arguing against the LS model, by drawing attention to plausible alternatives.

6. Accessibility-ranking and multiple activation

At this point, an advocate of the LS model may be tempted to downplay the contrast I have just stressed, between accessibility-based and literality-based models. The AS model and the LS model are in fact equivalent, she may argue, as it is not possible for some interpretation distinct from the literal interpretation to be more accessible than the latter. After all, non-literal interpretations on the picture I have sketched are accessed *through* literal interpretations. Now whatever activation is spread from the representation which constitutes the literal interpretation of the expression to other, associatively related representations cannot exceed the amount of activation present in the literal representation where the process of activation spreading originates. *It follows that a non-literal interpretation cannot be more accessible than the literal interpretation from which it proceeds*. Thus, even if we accept the AS model, according to which the most accessible interpretation is processed first, we need not reject the LS model, according to which the literal interpretation is processed first; for the literal interpretation *is* the most accessible one. If this is true we gain some generality; we can have a single serial model for all cases (ambiguity, sense construction, reference assignment, specifization and non-literal interpretation), namely the AS model, and consider the LS model as a particular version of the AS model (appropriate for cases where one candidate is distinguished as 'literal').

Although I find it highly desirable to have a single model for all cases, I strongly reject the conclusion that the literal candidate *is* the most accessible one. The reasoning which leads to that conclusion involves a fallacy very much like that which I exposed earlier, to the effect that only a serial model is available for non-literal discourse. The earlier reasoning used the asymmetric nature of non-literal meaning (its 'asymmetric dependence' on literal meaning) as evidence that the literal interpretation must be processed first. The fallacy lied in the fact that 'The literal interpretation is processed first' is

ambiguous: there is a distinction between processing the literal interpretation of the constituent (accessing it, as it were) and processing the literal interpretation of the *sentence* where it occurs. The former process is local, the latter global. The only thing which the asymmetric nature of non-literal interpretations entails is the priority of literal interpretations at the constituent level, but this is consistent with the absence of priority of literal interpretation at the sentence level. The parallel model and the AS model both incorporate this absence of priority.

The new fallacy, to the effect that there is no genuine distinction between the LS model and the AS model, also uses the asymmetric nature of non-literal interpretation as premiss. This asymmetry, in conjunction with the principle that a representation cannot transmit more activation than it possesses, is said to entail the following conclusion: that non-literal interpretations, insofar as they proceed from literal interpretations which spread activation to their neighbours, cannot be more accessible than the literal interpretations from which they get their activation. This is a fallacy, for, as Freud noticed long ago (Freud 1900, ch. 6), activation can be *over*-determined - it may come from different sources. Owing to this possibility of 'multiple activation', a representation may have a higher degree of activation (hence be more salient or accessible) than one of the representations from which it gets (part of) its activation, even if we accept that a representation cannot transmit more activation to associatively related representations than it possesses.

Let me say more precisely where I think the fallacy in the second literalist reasoning lies. The reasoning involves two basic principles as premises, namely the Conservative Principle (CP) and the Principle of Asymmetric Dependence (PAD):

(CP) A representation cannot transmit more activation than it possesses.
(PAD) Non-literal interpretations are accessed through literal interpretations.

The two principles together are said to entail that non-literal interpretations cannot be more accessible than literal interpretations (since they proceed from the latter, and a representation cannot transmit more activation than it possesses). But to get this result, we need something stronger than the PAD, and this is where the fallacy lies. In order to conclude (from the Conservative Principle) that non-literal interpretations cannot be more accessible than literal interpretations, we need what I call the *strengthened* PAD:

(Strengthened PAD):
Non-literal interpretations are accessed through, *and only through*, literal interpretations.

If non-literal interpretations are accessed through, and only through, literal interpretations, then whatever activation they have come from the literal interpretations; it follows that they can't be more accessible than the latter (in virtue of the Conservative Principle). But is it true that non-literal interpretations are accessed *only* through literal interpretations? No. The literal representation may be only one among many factors which contribute to the activation of the associatively related representation. As a result the latter may get a higher degree of activation than the literal representation itself possesses. It follows that the non-literal interpretation *can* be more accessible than the literal interpretation from which it proceeds, as a result of the simultaneous influence of *several* representations.

Consider the case of specifization I lused as example. When someone talks of 'wearing rabbit', the representation *rabbit fur* is activated as the result of (i) transmission of activation from the representation *rabbit* linguistically activated (*qua* literal semantic value of the term 'rabbit'), and (ii) transmission of activation from the associatively related representation *wearing* which is also linguistically activated. It is therefore theoretically possible for the representation *rabbit fur* to be more activated, in this context, than the more

general representation *rabbit*.[3] Whatever we think of this particular example, it seems to me that the following situation can arise: An expression linguistically encodes a certain representation; the latter is accessed when the expression is uttered, but another, associatively related representation is also activated (in part - but in part only - through the encoded representation) and turns out to be more salient in that context than the original representation from which (in part) it derives. On both the parallel model and the AS model, the derived representation goes directly into the overall interpretation of the utterance, without depending on the prior computation of the literal interpretation of the utterance involving the original representation. On the AS model *only* the derived representation goes into the overall interpretation of the utterance; on the parallel model both the derived representation and the literal representation go 'directly' into (distinct and parallel) interpretations of the utterance, without one interpretation being asymmetrically dependent on the other.

7. Accessibility shifts

Multiple activation need not be simultaneous. Some representation may be accessed through its links to the representation linguistically encoded and *become* more accessible than the latter as a result of the coming into play of further linguistic material. Consider sentence (1) again, or rather a variant of (1) in which the indefinite description 'a policeman' has been replaced by a name of the policeman to simplify matters:

(2) John was arrested by Mike yesterday; he had just stolen a wallet.

---

[3] This shows that one cannot consider (the mental representation) *rabbit fur* as being compositionally constructed out of the representations *rabbit* and *fur*. It has to be considered as an independent concept - perhaps a *de re* concept (Recanati 1993).

We may suppose that when the pronoun 'he' is uttered one of the two candidates, John or Mike, is more accessible than the other; John because he is the subject of the sentence, or Mike because he was mentioned last. For the sake of the argument, let us suppose that the second factor is more important than the first, in such a way that *Mike* is slightly more accessible than John when the pronoun is uttered. This may well change when the other constituents of the sentence are processed. When the predicate 'had just stolen a wallet' is uttered, *John* becomes more accessible than Mike as a candidate for the status of referent of 'he', even if Mike was more accessible at an earlier stage in the processing of the utterance. For John is the subject of 'was arrested' and therefore occupies the role of the person being arrested; now that role is linked to the role of the person doing the stealing, in some relevant frame (on frames, see e.g. Fillmore and Atkins 1992). Because of this link, the representation of the referent of 'he' as the person doing the stealing contributes some activation to the role of the person being arrested and therefore to John *qua* occupier of this role. John thus gains some extra activation and becomes the most accessible candidate.

John and Mike are on the same footing as far as literality is concerned (§5), but if we turn to a case in which some candidate is distinguished from the others by its being literal, we see that the same sort of temporal shift in accessibility can occur. When the words 'the ham sandwich' are uttered, we may consider that the representation of the ham sandwich is more accessible than other, related representations which are activated through their links to that representation. (This is a straightforward application of the CP.) Thus we may suppose that the representation of the ham sandwich is more accessible (more salient) than the 'derived' representation of the ham sandwich orderer. This is similar to the fact that Mike is more accessible than John when the pronoun 'he' in (2) is uttered, except that Mike is (perhaps) more accessible because he was mentioned last, while the representation of the

ham sandwich is more accessible because it is linguistically encoded.[4] In both cases, the initial ranking is reversed when further linguistic material comes into play. After the predicate in the sentence 'The ham sandwich has left without paying' has been processed, the ham sandwich is no longer a more accessible candidate than the ham sandwich orderer - the order of accessibility is reversed. The explanation, again, is very simple and does not appeal to inference on the hearer's part. The predicate 'has left without paying' demands a person as argument; this raises the accessibility of all candidates who are (represented as) persons. In this way the ham sandwich orderer gains some extra activation which makes him more accessible than the ham sandwich, after the predicate has been processed.

Accessibility shifts of the sort I have just described will turn out to be important for explaining garden-path effects within an accessibility-based framework (see §10 below). The notion of an accessibility shift also provides immediate answers to the questions which, according to Kent Bach (forthcoming), my account of the ham sandwich metonymy and similar phenomena in terms of local pragmatic processes raises. Bach writes:

Recanati supposes that the process of metonymical transfer takes place without the intrusion of a thought of the absurd proposition associated with the literal meaning of (i) ['The ham sandwich is getting restless']. That is, the hearer does not have to compute that the speaker does not mean that a certain culinary item is getting restless in order to determine what the speaker does mean. But, I ask, how can the hearer go from the concept of ham sandwich to that of ham-sandwich-orderer without first entertaining the absurd minimal proposition? What triggers the 'local process' and, for that matter, keeps it from being triggered in a cases like (ii) ['The ham sandwich is getting eaten'], uttered in similar circumstances? Recanati's account predicts that the hearer would entertain the proposition that the ham-sandwich-orderer is getting

---

[4] For qualifications, see fn. 2.

eaten, since the local process it posits would get triggered before a full proposition were reached. And yet (ii) could be understood perfectly well. (Bach forthcoming, §8)

But my account does *not* predict that the hearer would understand (ii) ('The ham sandwich is being eaten') as being about the ham sandwich orderer, for the reason that there *is* something, in my account, which keeps the local process of metonymy from being triggered in a case like (ii). What keeps the local process of metonymy from being triggered in a case like (ii) is precisely what answers Bach's legitimate question: 'How can the hearer go from the concept of ham sandwich to that of ham-sandwich-orderer without first entertaining the absurd minimal [i.e. literal] proposition?' In my framework, the process of metonymical transfer is 'local', not global, yet it is sensitive to the linguistic (and extralinguistic) context in which the expression which receives the metonymical interpretation occurs. In particular, it is sensitive to the meaning of the predicate expression; this is why it is triggered in (i) but not in (ii).

For the literal proposition to be entertained, the literal semantic value of 'the ham sandwich' must undergo semantic composition with the semantic value of the predicate; but the semantic value of the predicate is a concept whose very activation makes the metonymical interpretation of the subject-term more accessible than its literal interpretation, hence prevents the latter from going into the proposition expressed. The activation of the concept which is the literal interpretation of the predicate is both what triggers the process of metonymical transfer in (i) and what prevents it from taking place in (ii). Yet the transfer's sensitivity to the interpretation of the other constituents in the sentence does not mean that the transfer is a global process presupposing the computation of the absurd, literal proposition. The interpreter does not go from the concept of ham sandwich to that of ham-sandwich-orderer after having entertained the absurd literal proposition; rather, it is *because* the interpreter goes from the concept of ham sandwich to

24

that of ham-sandwich-orderer (as a result of an accessibility shift resulting from the interpretation of the predicate) that he or she does *not* entertain the absurd literal proposition.

In this account, the metonymical transfer from ham sandwich to ham-sandwich-orderer results from a shift in accessibility triggered by the interpretation of another constituent in the sentence. The literal interpretation of the subject-term was more accessible than the metonymical interpretation before the predicate came into the picture, but the metonymical interpretation becomes more accessible as a result of interpreting the predicate. This accessibility shift has two consequences. First, the literal value of the subject term in (i) does not undergo semantic composition with that of the predicate, precisely because the semantic value of the predicate makes the literal value of the subject term less accessible than some other, non-literal interpretation. Second, if the sentence contained *another* predicate which did *not* require a person, but rather a culinary item, as argument, then the accessibility shift would *not* occur and the description 'the ham sandwich' would be given the literal interpretation; this account for the difference between (i) and (ii).

8. Interactive processing: the role of schemata

So far I have considered the interpretation of a given constituent as if the interpretations of the other constituents were fixed; but nothing is fixed. All constituents can be interpreted non-literally as well as literally, and how we interpret a constituent cannot but affect how we interpret the others. Consider (3), for instance.

(3)   The city is asleep.

A city is not the sort of thing that sleeps (in the normal sense of the term), hence the overall interpretation of (3) is likely to involve some process of non-literal interpretation. There are various possibilities, though. 'The city' can be interpreted non-literally (metonymically) as standing for 'the inhabitants of the city'; or 'asleep' can be interpreted metaphorically in the sense of 'quiet and showing little activity'. Given the likelihood of a non-literal interpretation for some constituent, if 'the city' literally applies to a city, 'asleep' will have to be taken non-literally; conversely, if 'asleep' is literal, 'the city' will not   be.

The same sort of trade off applies to seemingly unproblematic examples such as the following:

(4) I finished the book.
(5) John heard the piano.

If 'finish' has its standard sense in (4), its object (what is said to be finished) must be a process, for only processes can start or finish in the standard sense (Pustejovski 1991). It follows that 'the book' must be interpreted non-literally - it must be interpreted as standing, not for a certain book, but for the process of, say, *reading* a certain book. Conversely, if 'the book' is interpreted literally and stands for a certain book, 'finish' must have a derived sense and mean something like *finished reading*.[5] Pustejovski (1991) defends the former analysis, Langacker (1984) the latter.[6]

---

[5]In the latter case, 'finish' first acquire a derived meaning *finish X-ing* involving a variable and the latter is contextually instantiated, in such a way that the sense *finish reading* results from a two-step process of contextual interpretation.

[6]The sentence may also be seen as elliptical and contextually 'expanded' into 'I finished reading the book' (Bach 1987, forthcoming). I find this sort of analysis too much unconstrained (Recanati 1993:261ff), so I don't take it into consideration here.

(5) is another example (discussed by Langacker) which can be interpreted in two ways. We may construe 'the piano' as standing for *the sounds emitted by the piano* (metonymy), in accordance with the 'postulate' that only sounds can be heard; or we may consider 'hear' as polysemous. When it is said that 'only sounds can be heard', 'hear' is taken in its basic sense ('hear'$_1$), but there is another sense ('hear'$_2$), which can be defined as follows: An object is heard$_2$ whenever the sound it emits is heard$_1$. Langacker analyses (5) as saying that John heard$_2$ the piano, rather than as saying that he heard$_1$ the sounds emitted by the piano.

Though (3), (4) and (5) exhibit the same sort of trade off, there is an important difference between (3), on the one hand, and (4) and (5) on the other hand. Depending on which constituent receives a non-literal interpretation, the global interpretation of (3) itself will vary. In one case the speaker means that the inhabitants of the city are sleeping, in the other she means that the city itself is quiet and shows little activity; the truth-conditions are distinct. It follows that it is easy to say which constituent is responsible for the non-literal interpretation of the sentence: we have only to consult our intuitions concerning the truth-conditions of the utterance. (4) and (5) are very different, however. Whether one constituent or the other bears the onus of non-literal interpretation, the overall interpretation of the utterance does not seem to change. Hence it is not easy to establish which analysis is right, and one must appeal to indirect evidence, such as, for example, the fact that we can say things like: 'I can both hear and touch the piano'. (This fact shows that Langacker is right: when we say that we hear the piano, we are referring to the piano. It is the verb, not the noun-phrase, which has a derived sense in this context.)[7]

---

[7]A number of similar tests are offered in Nunberg 1993 and forthcoming. For example, Nunberg argues that (a) and (b) must be analysed in fairly different ways, despite their superficial similarity.

(a) I am parked out back.

The trade off talked about in this section is a particular case of a more general and fairly complex phenomenon: the search for coherence in interpretation. We have two constituents, A and B, each with (at least) two possible interpretations (A_i and A_j, and B_i and B_j). Now there is a better 'fit' between A_i and B_j or between A_j and B_i than there is between either A_i and B_i or A_j and B_j, hence the 'trade-off' effect: the two global interpretations $<A_i, B_j>$ and $<A_j, B_i>$ are favored, in such a way that the choice of either A_j or A_i as the correct interpretation for A 'coerces' a particular interpretation (B_i or B_j) for B.The tendency to prefer coherent interpretations (with a high degree of fit between the

---

(b) This one <pointing to a car key> is parked out back.

(b) is a case of 'deferred ostension': an object is 'demonstrated' (the key), but the ultimate 'referent', as opposed to the 'demonstratum', is another object suitably related to the demonstratum, namely the car itself. In contrast, 'I' does *not* refer to the car in (a); it normally refers to the speaker. It's the *predicate* 'parked out back' that bears the onus of non-literality. That the referent is the car rather than the (demonstrated) key in (b) is established by the fact that semantic features of the referring expression such as number or gender do NOT correspond to the key, but to the car; thus we would say 'this one' (in the singular) even if we were exhibiting a whole set of keys for the car in   question. Also, we can say things like:

(c) This <pointing to a car key>, which is an old Chevrolet, is parked out back.

The relative clause clearly qualifies the car, not the key. These tests provide evidence that the referent is the car, even if it is the key which is 'demonstrated'. When we turn to (a), the situation is totally different. Gender and number correspond to the speaker, not to the car. We would not say 'We are parked out back' if there was a single car owner and several cars (whereas we can say ' those <pointing to a key> are parked out back' if there is a single key and several cars). And we can't say   'I, who am an old Chevrolet, am parked out back'.

various semantic values) is what we must now try to account for. But first, we must introduce the notion of a schema (Rumelhart 1979, 1980; Langacker 1987).

To say that two semantic values α and β 'fit together' is to say that there is an abstract schema which <α, β> instantiate. The particular case discussed above, 'coercion', can be analysed as follows. An expression E activates an abstract schema in which there is a slot for a value of a certain type; as a result, the semantic value of E will preferably enter into composition with a semantic value of the relevant type. Thus in a sentence like 'the city is asleep', if we give to 'asleep' its literal value (thereby activating the SLEEP schema), the value of 'the city' will have to be of the relevant type (e.g. human or animal), hence non-literal.

The role of schemata in interpretation is best seen in connection with more complex examples. Consider example (2) from section 7:

(2) John was arrested by Mike yesterday; he had just stolen a wallet.

As I said earlier, John is selected as the referent for 'he' in the second clause because (i) the referent of 'he' is said to have stolen and John is known to have been arrested, and (ii) there is a frame or schema in which the two roles (stealing, and being arrested) are linked. This schema is jointly activated by the predicates 'was arrested' and 'had stolen'. An interpretation in which the same person steals and is arrested (and in which he is arrested because he has stolen - see below) satisfies the schema, and is more likely to be selected than one which violates it.

Given the link we have established between coherence (or fit) and schemata, the question we must answer becomes: Why are schema-instantiating interpretations more successful than others? In line with what has been said so far, the following answer suggests itself. Interpretational success - what brings a 'candidate' or potential semantic value into the actual interpretation of the utterance - to a large extent depends on the

candidate's accessibility or degree of activation. Now a schema is activated by, or accessed through, an expression whose semantic value corresponds to an aspect of the schema. The schema thus activated in turn raises the activation level (the accessibility) of whatever possible semantic values for other constituents of the sentence happen to fit the schema. The schema itself - hence the semantic value through which it is accessed - gains extra activation from the fact that some other constituent of the sentence has a possible interpretation which fits the schema. In such a case all 'candidates' or potential semantic values which fit the schema evoked by some of them mutually reinforce their accessibility and therefore increase the likelihood that they will be globally selected as part of the interpretation of the utterance. Coherent, schema-instantiating interpretations therefore tend to be selected and preferred over non-integrated or 'loose' interpretations. As a result, schemata drive the interpretation process.

The role played by schemata explains why the process of utterance interpretation is to such a large extent top down and driven by world knowledge (Rumelhart 1979; Travis 1975, 1981; Searle, 1978, 1980). The interpreter unconsciously enriches the situation described by the utterance with many details which do not correspond to any aspect of the uttered sentence but are contributed in order to fit an evoked schema (on 'enrichment', see Recanati 1993 and the references therein).[8] Thus, in the John-Mike example above, not

---

[8] Think of an example like (6):

(6) John hates the piano.

Contrary to (3)-(5), this example cannot be described in terms of 'coercion'. I mentioned earlier various contentions to the effect that only sounds can be heard and only processes finished. By virtue of these principles, some process of non-literal interpretation must take place in order to make sense of phrases such as 'finish the book' or 'hear the piano'. But a piano is certainly an object that *can* be hated, however strictly one construes the predicate 'hate'. Still, some contextual enrichment is in order, because to hate the piano is to hate it

only is the reference of 'he' determined by the *Steal (x) - Is arrested (x)* schema, but (among other things) the causal interpretation of the relation between the two events mentioned - the fact that the referent was arrested because he had stolen - is also determined by the schema. Specifization and a good deal of what I earlier called sense construction are similar cases of schema-driven enrichment (Clark 1992, Goschke & Koppelberg 1990, Langacker 1987, Medin & Shoben 1988, Murphy 1988 and 1990, Ruhl 1989, Nunberg & Zaenen forthcoming).

9. When is an interpretation satisfactory?

What has just been said about schemata and accessibility enables us to complete the picture of utterance interpretation sketched in the earlier sections of this paper. When discussing the parallel and serial models, I made an implicit distinction between two mechanisms: the mechanism that determines which candidate semantic value for a given constituent is 'tried' and undergoes composition with other semantic values corresponding to other constituents of the sentence; and the mechanism which determines whether or not an interpretation which has been tried is actually retained. I discussed three possible models or classes of models for the first mechanism, but remained fairly vague as to the nature of the second mechanism; I only said that an interpretation is retained if it 'fits the broader context of discourse'. (If a single interpretation is tried, as in the two serial models, that interpretation is retained if it fits the broader context of discourse and rejected otherwise; if several interpretations are tried, as in the parallel model, the interpretation to be retained is that which best fits the broader context of discourse.)

---

under some aspect or dimension. One may hate the sounds emitted by the piano, or one can hate playing the piano, or one can hate the piano as a piece of furniture, etc. The relevant dimension is contextually provided through the process of enrichment.

What is it for an interpretation to 'fit the broader context of discourse'? I think we are now in a better position to answer this question. In relating an utterance to its broader context, we look for coherence, i.e. schema-satisfaction, as much as we do when interpreting a given constituent in the context of the utterance where it occurs. The same principles apply at both levels (as a simple two-sentence example like (2) already shows). It follows that the second mechanism is not significantly different from the first, within an accessibility-based framework. Accessibility not only determines which candidate is considered and processed, it also determines which candidate is retained.

Dan Sperber objected to me in conversation that accessibility can lead us astray: Sometimes the first interpretation that comes to mind (the most accessible one) turns out not to be satisfactory and forces the hearer to backtrack. According to Sperber, the possibility of such garden-path effects shows that interpretational success cannot be equated with sheer accessibility. This objection is misguided, I think. The most accessible interpretation at some stage s in the interpretation process may well turn out to be unsatisfactory at some later stage s', thereby resulting in a garden path effect and the need to backtrack. This does not show that interpretational success cannot be cashed out in terms of accessibility. At any given stage, the most accessible interpretation will be the winning one (at that stage). In garden path utterances we have *two* successive stages to consider. Some interpretation is the most accessible one, hence wins, at s, but that interpretation fails to fit some schema, hence loses, at a later stage s'. In an accessibility-based framework, this means that this interpretation's accessibility at s' is no longer sufficient for it to be the winning candidate (at s'). Another candidate (which was less accessible at s, but turns out to be more accessible at s') takes over, hence the garden path effect. The distinction between successive stages of interpretation, together with the notion of an accessibility shift, is sufficient to account for garden path effects within the accessibility-based framework.

10. Counterexamples to the LS model

So far I have tried to make a case for an accessibility-based alternative to the LS model. To me at least the accessibility-based model seems much more plausible than the LS model. In this last section before concluding I want to go further and provide a direct counter-example to the LS model.

The sort of counter-example to the LS model which I want to construct exploits the interactive nature of interpretation. As we saw in §8, how we interpret a particular constituent depends upon how we interpret the other constituents. The cases described in §8 are cases in which one constituent has to be interpreted non-literally if the others are interpreted literally. There are also cases in which a non-literal interpretation for a constituent prompts a non-literal interpretation for another constituent, and cases of interaction involving other types of contextual interpretation (e.g. sense selection or sense construction). Thus in interpreting an utterance of: 'Mike finished John's book', we have to assign a contextual value to 'finish' (*finish reading*, *finish writing*, *finish binding*, etc.) and a contextual value to 'John's book' (*book owned by John*, *written by John*, *read by John*, etc.). There is an obvious interaction between the two values respectively assigned to the two constituents; the overall assignment must be coherent. If 'John's book' is interpreted as meaning *the book written by John,* it's unlikely that 'finished' will be interpreted as meaning *finished writing*.

The cases of interaction I am interested in here are cases in which the contextual processes which interact respectively belong to the two broad classes distinguished in my book *Direct Reference* (Recanati 1993: ch. 13-14). The two classes in question are: the contextual processes which are semantically *mandatory* because no proposition would be expressed without them, and the contextual processes which are *optional*, i.e. non-mandatory. Optional processes take a complete proposition as input and yield another proposition as output. Enrichment and non-literal interpretation, in most cases at least, are

optional processes. Thus we normally interpret 'He took out his key and opened the door' as meaning that he opened the door with the key, but even if we did not 'enrich' the interpretation of the utterance in this way the latter would still express a complete proposition (namely the proposition that he did two things: take out his key, and open the door in some way or other). In this regard enrichment differs from the sort of contextual process involved in sense construction; for 'I read John's book' does not express a complete proposition if a certain relation between John and the book is not contextually provided.

It is possible for two contextual processes respectively belonging to the two classes distinguished above to interact in such a way that *the mandatory process depends upon the optional one*, and cannot operate until the latter has. For example, take the construction: the x of y, or even better the possessive construction: y's x. For the sentence where such a construction occurs to express a definite proposition, a certain relation between x and y must be contextually provided. Now which relation is actually provided as part of the interpretation of the utterance obviously depends upon the identity of the relata and what we know about them. The relata, x and y, will themselves be referred to by means of linguistic expressions, say NP_1 and NP_2. Now it's easy to imagine a case in which, say, NP_2 can undergo a process of non-literal interpretation. For example a definite description such as 'the lion' can be either interpreted literally (as referring to a lion) or non-literally (as referring metonymically to, say, the warrior whose shield carries a picture of a lion). How we interpret the description determines the identity of one of the relata, and we have seen that (in some cases at least), how we interpret the relation expressed by the possessive construction depends upon the identity of the relata. I want to argue that, in such a case, if the description is used non-literally (if it refers to the warrior), then a contextual process which must take place for the utterance to express a complete proposition (namely the assignment to the possessive construction of a particular relation as contextual value) cannot take place unless a non-literal interpretation has been antecedently provided for some other constituent in the sentence. It follows that the utterance is *semantically*

indeterminate at the purely literal level, for the mandatory process of sense construction - without which no complete proposition would be expressed - itself depends on some feature of the *non-literal* interpretation of the utterance and therefore cannot be achieved at the literal level. This is a case in which it is not possible (even in principle) to determine the proposition literally expressed before one has determined what the speaker non-literally means.

Let us now actually construct the example by filling in the context. Suppose that there has been a ritual fight between respectively five warriors and five beasts. The beasts are a wolf, a lion, a hyena, a bear, and an alligator; the warriors are armed with swords, and carry shields with distinctive decorations (the first warrior has the moon on his shield, the second one has the Eiffel Tower, the third one has the Metro-Godwin-Mayer roaring lion, etc.). Each warrior is assigned a particular beast which he or she must stab to death. After the ritual fight, the five beasts lay on the ground with a sword through their body. This is the context. Now suppose that in this context I utter (7):

(7)   Bring me the lion's sword - I want to have a look at it.

In this context I think there are two accessible interpretations for 'the lion'; the lion can be interpreted literally as referring to the lion (one of the beasts), or non-literally as referring to the warrior who has (a picture of) a lion on his shield. If we choose the first interpretation, the relation which will be contextually provided will be one of the salient relations which hold between the lion (the animal) and the sword which can be seen emerging from its pierced body. If we choose the second interpretation, the relation will be totally different; it will be one of the salient relations between the warrior and the sword which he used in his fight against, say, the bear.

Now suppose the correct interpretation is the second one: the speaker actually refers to the warrior with a lion on his shield and wants to see his sword (i.e. the sword which

emerges from the bear's body). My contention is that the LS model cannot account for a case like this. For it's not possible, even in principle, to compute the proposition literally expressed before accessing the non-literal interpretation, in a case like this.

11. Conclusion

In this paper I have argued against a widespread model of utterance interpretation, namely the LS model, according to which the literal interpretation of an utterance (the proposition literally expressed by that utterance) must be computed before non-literal interpretations can be entertained. I have tried to make room for the possibility that an utterance receives a non-literal interpretation (i) as a result of a *local* process concerning only a particular constituent of the sentence, and (ii) without the literal interpretation *of the utterance* necessarily being computed. On my view, the literal interpretation of the constituent must be accessed, but there is no need to take it into the global interpretation of the utterance; the latter may recruit only the non-literal interpretation of the constituent, in such a way that the literal interpretation of the global utterance is not computed, or at least need not be computed for the non-literal interpretation to be entertained. (It may be computed in parallel to the non-literal interpretation, but at least the latter is not asymmetrically dependent on the former, contrary to what the LS model claims.)

Two models were presented as alternatives to the LS model. On the parallel model, literal and non-literal interpretations of the utterance are simultaneously processed. On the AS model, only the most accessible candidate for the status of semantic value of the relevant constituent goes into the overall interpretation of the utterance; if the most accessible candidate is not the literal semantic value of the constituent, the literal interpretation of the sentence is not computed - only the non-literal interpretation of the sentence (involving the highly accessible but non-literal interpretation of the constituent) is delivered.

36

I have considered two arguments in favour of the LS model. The first argument purports to show that the parallel model is ruled out in principle, because of the asymmetric nature of non-literality (the fact that non-literal meaning asymmetrically depends on literal meaning). The second argument also appeals to the asymmetric nature of non-literality to rule out the AS model construed as an alternative to the LS model. Both arguments have been shown to be fallacious: the asymmetric dependence of non-literal meaning on literal meaning does *not* support the LS model as opposed to the parallel model or the AS model. If I am right the LS model is baseless: there is no particular reason why the process of non-literal interpretation should proceed serially through computing the literal interpretation of the sentence. At the end of the paper I have gone one step further. I have provided data which, it seems to me, the LS model cannot account for. I conclude that the LS model is not only unwarranted but also empirically inadequate.

References

J. Atlas (1989). *Philosophy without Ambiguity: A Logico-Linguistic Essay.* Oxford: Clarendon Press.

K. Bach (forthcoming). Conversational Impliciture. To appear in *Mind and Language*.

L.Barsalou & D. Billman (1989). Systematicity and Semantic Ambiguity. In D. Gorfein (ed.), *Resolving semantic ambiguity*. New York: Springer Verlag.

K. Bühler (1934). *Sprachtheorie*. Jena: Gustav Fischer [Engl. trans. by D. F. Goodwin, *Theory of Language*. Amsterdam: John Benjamins, 1990].

R. Carston (1988). Implicature, Explicature, and Truth-Theoretic Semantics. In R. Kempson (ed.), *Mental Representations: the Interface between Language and Reality*. Cambridge: Cambridge University Press.

H. Clark (1979). Responding to Indirect Speech Acts. *Cognitive Psychology* 11, 430-77.

H. Clark (1992). *Arenas of Language Use*. Chicago: The University of Chicago Press and Center for the Study of Language and Information.

L. J. Cohen (1985). How is Conceptual Innovation Possible. *Erkenntnis* 25, 221-38.

M. Dascal (1987). Defending Literal Meaning. *Cognitive Science* 11, 259-81.

C. Fillmore & B. Atkins (1992). Toward a Frame-Based Lexicon: the Semantics of RISK and its Neighbors. In A. Lehrer & E. F. Kittay (eds.), *Frames, Fields and Contrasts: New Essays in Semantic and Lexical Organization*. Hillsdale: Lawrence Erlbaum Associates.

S. Freud (1900). *Die Traumdeutung*. Leipzig: F. Deuticke.

M. Garman (1990). *Psycholinguistics*. Cambridge: Cambridge University Press.

R. Gerrig & A. Healy (1983). Dual Processes in Metaphor Understanding: Comprehension and Appreciation. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 9, 667-675.

R. Gerrig & G. Murphy (1992). Contextual Influences on the Comprehension of Complex Concepts. *Language and Cognitive Processes,* 7, 205-30.

R. Gibbs (1983). Do People Always Process the Literal Meanings of Indirect Requests? *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 9, 524-533.

R. Gibbs (1984). Literal Meaning and Psychological Theory. *Cognitive Science* 8, 275-304.

R. Gibbs (1986a). On the Psycholinguistics of Sarcasm. *Journal of Experimental Psychology: General,* 115, 3-25.

R. Gibbs (1986b). Skating on Thin Ice: Literal Meaning and Understanding Idioms in Conversation. *Discourse Processes* 9, 17-30.

R. Gibbs & R. Gerrig (1989). How Context Makes Metaphor Comprehension Seem "Special". *Metaphor and Symbolic Activity* 4, 145-58.

P. Gildea & S. Glucksberg (1983). On Understanding Metaphor: the Role of Context. *Journal of Verbal Learning and Verbal Behavior* 21, 512-21.

T. Goschke & D. Koppelberg (1990). Connectionist Representations, Semantic Compositionality, and the Instability of Concept Structure. *Psychological Research* 52, 253-70.

P. Grice (1989). *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.

R. Harris (1976). Comprehension of Metaphors: a Test of the Two-Stage Processing Model. *Bulletin of the Psychonomics Society* 8, 312-14.

A. Inhoff, S. Lima & P. Carroll (1984). Contextual Effects on Metaphor Comprehension in Reading. *Memory and Cognition* 12, 558-67.

D. Kaplan (1989). Demonstratives. In J. Almog, H. Wettstein and J. Perry (eds.), *Themes from Kaplan*. New York: Oxford University Press.

P. Kay & K. Zimmer (1976). On the Semantics of Compounds and Genitives in English. *Sixth California Linguistics Association Conference Proceedings*, San Diego: Campanile.

R. Kempson (1977). *Semantic Theory*. Cambridge: Cambridge University Press.

R. Langacker (1984). Active Zones. *Proceedings of the Annual Meeting of the Berkeley Linguistics Society* 10: 172-88.

R. Langacker (1987). *Foundations of Cognitive Grammar*. Stanford: Stanford University Press.

D. Medin & E.J. Shoben (1988). Context and Structure in Conceptual Combination. *Cognitive Psychology* 20, 158-90.

G. Murphy (1988). The Comprehension of Complex Concepts. *Cognitive Science* 12, 529-62.

G. Murphy (1990). Noun Phrase Interpretation and Conceptual Combination. *Journal of Memory and Language* 29, 259-88.

G. Nunberg (1979). The Non-Uniqueness of Semantic Solutions: Polysemy. *Linguistics and Philosophy* 3, 143-84.

G. Nunberg (1993). Indexicality and Deixis. *Linguistics and Philosophy* 16, 1-43.

G. Nunberg (forthcoming). Transfers of Meaning. To appear in *Journal of Semantics*.

G. Nunberg & A. Zaenen (forthcoming). Systematic Polysemy in Lexicology and Lexicography.

A. Ortony, D. Schallert, R. Reynolds & S. Antos (1978). Interpreting Metaphors and Idioms: Some Effects of Context on Comprehension. *Journal of Verbal Learning and Verbal Behavior* 17, 465-77.

B. Partee (1984). Compositionality. In F. Landman & F. Veltman (eds.), *Varieties of Formal Semantics*. Dordrecht: Foris.

J. Pustejovsky (1991). The Generative Lexicon. *Computational Linguistics* 17, 409-41

F. Recanati (1989). The Pragmatics of What is Said. *Mind and Language* 4, 295-329 [reprinted in S. Davis (ed.), *Pragmatics: A Reader*. Oxford University Press 1991].

F. Recanati (1993). *Direct Reference: From Language to Thought*. Oxford: Basil Blackwell.

C. Ruhl (1989). *On Monosemy: A Study in Linguistic Semantics*. Albany: State University of New York.

D. Rumelhart (1979). Some Problems with the Notion of Literal Meaning. In A. Ortony (ed.), *Metaphor and Thought*. Cambridge: Cambridge University Press.

D. Rumelhart (1980). Schemata: the Building Blocks of Cognition. In R. Spiro, B. Bruce, and W. Brewer (eds.), *Theoretical Issues in Reading Comprehension*. Hillsdale, N.J: Lawrence Erlbaum.

I. Sag (1981). Formal Semantics and Extralinguistic Context. In P. Cole (ed.), *Radical Pragmatics*. New York: Academic Press.

J. Searle (1978). Literal Meaning. *Erkenntnis* 13, 207-24.

J. Searle (1980). The Background of Meaning. In J. Searle, F. Kiefer, and M. Bierwisch (eds.), *Speech Act Theory and Pragmatics*. Dordrecht: Reidel.

D. Sperber & D. Wilson (1986). *Relevance: Communication and Cognition*. Oxford: Basil Blackwell.

C. Travis (1975). *Saying and Understanding*. Oxford: Basil Blackwell.

C. Travis (1981). *The True and the False: the Domain of the Pragmatic*. Amsterdam: John Benjamins.