



WILEY

---

Loose Talk

Author(s): Dan Sperber and Deirdre Wilson

Source: *Proceedings of the Aristotelian Society*, New Series, Vol. 86 (1985 - 1986), pp. 153-171

Published by: [Wiley](#) on behalf of [Aristotelian Society](#)

Stable URL: <http://www.jstor.org/stable/4545042>

Accessed: 13-01-2016 14:24 UTC

---

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



*Aristotelian Society and Wiley* are collaborating with JSTOR to digitize, preserve and extend access to *Proceedings of the Aristotelian Society*.

<http://www.jstor.org>

## IX\*—LOOSE TALK

by Dan Sperber and Deirdre Wilson

Literal talk, loose talk and metaphorical talk are often seen as different in kind. We want to argue that they differ not in kind but only in degree of looseness, and that they are understood in essentially the same way. The literature on literalness, looseness (or vagueness) and metaphor is vast; we will not attempt to review it here. Our discussion, which will be both brief and untechnical, is based on a view of human communication developed in greater detail in our book *Relevance: Communication and cognition* (Sperber and Wilson 1986).

### I

#### *The issue*

Suppose Mary believes that the car is in the garage and intends Peter to share this belief. One way she can fulfil her intention is by informing him of it: knowing her intention, he will have good reason for fulfilling it provided that he trusts her. One way she can inform Peter of her intention is by saying to him:

(1) The car is in the garage.

This utterance expresses just the proposition and propositional attitude that Mary intends Peter to share with her.

As it stands, this commonplace account leaves several questions unanswered. In particular, it does not explain why Peter should take utterance (1) as evidence that Mary has, and intends him to adopt, the propositional attitude expressed. One generally accepted answer is that there is a rule (or norm, or principle, or maxim, or convention, or presumption) of literal truthfulness whereby the utterer of a declarative sentence, in expressing a certain proposition, automatically vouches for its truth (similar rules of literal commitment can be formulated for non-declarative utterances). This accounts well enough for examples such as (1). It is tempting to go on to treat these

\* Meeting of the Aristotelian Society held at 5/7 Tavistock Place, London WC1, on Monday, 3 March, 1986 at 6.00 p.m.

examples as paradigmatic of verbal communication in general.

However, there are various exceptions to a hypothetical rule of literal truthfulness. Some exceptions (e.g. quotations) involve declarative sentences whose literal meaning is conveyed, but not asserted. We will not consider these today. Others involve a departure from literalness: the speaker is apparently vouching for the truth of *some* proposition, but not for the truth of the proposition literally expressed. The most blatant exceptions of this second type are metaphors such as (2):

(2) *Mother to child*: You're a piglet.

The mother who says to her child that he is a piglet is certainly not committed to the literal truth of her utterance. She seems, rather, to be vouching for the truth of some proposition such as (3):

(3) He is a dirty child.

Metaphors, and tropes in general, are classically described as departures from a norm of literalness: a 'figurative meaning' such as (3) is said to be substituted for the literal meaning of (2). But, as romantic critics of classical rhetoric have pointed out, ordinary discourse is shot through with metaphors; if anything, it is a long stretch of strictly literal discourse that should be seen as a departure from a norm. The initial implausibility of any hypothetical rule of literal truthfulness might be overlooked if the appeal to such a rule had useful theoretical consequences; if it helped to explain how not only literal talk, but also loose talk and metaphor are understood. But in this respect, modern accounts are neither essentially different from, nor superior to, classical rhetorical accounts.

For instance, Grice's brief account of figurative language (Grice 1975:53; Grice 1978:123-125) is very much in the classical rhetorical tradition. He treats irony, metaphor, hyperbole and litotes as departures from a norm, the norm in this case being obedience to a 'maxim' of truthfulness. According to Grice, when a speaker says, or makes as if to say, something which would blatantly violate the maxim of truthfulness, the hearer will assume that the maxim is being observed on another level, and will try to recover as an implicature some related proposition which a speaker observing the maxims

might have wanted to convey. For instance, in example (2)-(3) above, the child would infer from the fact that his mother could not have intended to assert truthfully and literally that he is a piglet, that she must have implicated the related proposition that he is a dirty child.

What is classically treated as a figurative meaning is thus reanalysed by Grice as an implicature. Both approaches assume that when the literal interpretation is inappropriate, the appropriate figurative interpretation somehow comes to the hearer's mind. Both resort, implicitly or explicitly, to some form of associationist psychology: trains of thought are seen as guided by contiguity, part-whole relations, resemblance and antinomy. Such views are no longer considered adequate to account for other cognitive abilities, but are still called upon, for want of any alternative, when it comes to explaining what is evoked by a metonymy, a synecdoche, a metaphor, or an irony. No other explanation is given of how figurative interpretations are recovered. Grice's account merely adds an inferential step of confirmation to these mysteriously retrieved figurative interpretations. Yet by Grice's own criteria, an implicature should be not only confirmable, but also calculable.

Many modern pragmatic accounts, including Grice's, are vulnerable to another criticism raised by the romantics against classical rhetoricians. The romantics challenged the classical view, shared by many modern pragmatists, that any trope has a literal paraphrase which is cognitively synonymous. Against this, the romantics maintained that a felicitous trope cannot be paraphrased. Thus Coleridge argues that the 'infallible test of a blameless style' is:

its *untranslateableness* in words of the same language without injury to the meaning. Be it observed, however, that I include in the *meaning* of a word not only its correspondent object but likewise all the associations which it recalls.  
(Coleridge: *Biographia Literaria*, Ch. XXII)

In her modest way, the mother who calls her child a piglet achieves some unparaphrasable effects: for instance, she seems more indulgent than if she had called him a dirty child. More generally, utterances in which all the speaker wants to do is to inform the hearer of a simple fact are untypical of communication

in general. Quite often, the speaker wants to communicate not a single atomic proposition, but a complex thought made up of many atomic thoughts, some of which are salient while others are not consciously spelled out in her mind. The speaker does not expect the hearer to entertain exactly the same complex thought. Rather, she intends him to entertain the proposition(s) most salient in her mind and to construct around it (or them) a complex thought which merely bears some similarity to her own. For instance, the mother wants the child to realise quite clearly that she thinks he is dirty, and to get at least an inkling of her accompanying thoughts. Some implicatures, we will claim, are strongly conveyed, others are weakly conveyed: implicatures come in varying degrees of strength.

The romantic critics were unquestionably right to draw attention to the richness and importance of those effects of figures of speech which are not maintained under paraphrase. These effects have merely been noted by classical rhetoricians and modern pragmatists alike; they have not been described, let alone explained, and have been treated without further discussion as cognitively negligible ornaments. But for all their justified criticisms and subtle observations, the romantics and their modern descendants have been content to talk about metaphor in metaphorical terms, and have proposed no explicit theory of their own; if anything, they have cast doubt on the very possibility of a non-metaphorical theory of metaphor by rejecting outright both the notion of a literal meaning—the ‘proper meaning superstition’ as I.A. Richards calls it—and the framework of truth-conditional semantics.

Our aim here is to give a brief sketch of a theory which differs from both the classical and romantic approaches and from their modern counterparts. Unlike romantic theorists, we will accept the idea that utterances express a literal truth-conditional meaning which is partly determined by the semantics of the sentence uttered; unlike classical theorists, we will challenge the idea that the speaker normally communicates the literal meaning of her utterance.

## II

### *Resemblance*

An utterance expresses a proposition. In consequence, it

represents a state of affairs: the state of affairs which must obtain for the proposition expressed to be true. We would like to suggest, though, that utterances are not restricted to representing states of affairs. Any object in the world can, in principle, be used to represent any other object that it resembles. For instance, a piece of rope can be used to represent a snake which it resembles in shape. An utterance can be used to represent another utterance which it resembles in meaning—either closely, as in the case of a paraphrase or translation, or more distantly, as in the case of a summary. Generally speaking, an utterance can be used to represent any representation which it resembles in content, whether a public representation such as another utterance, or a mental representation such as a thought.

To distinguish these two modes of representation—representation in virtue of truth-conditions and representation in virtue of resemblance—we will call the former *description*, and the latter *interpretation*. We will say that an utterance *descriptively* represents the state of affairs which makes the proposition it expresses true, and *interpretively* represents a representation which it resembles in content. While resemblance in general is a notoriously vague notion, we are only interested here in resemblance in content between representations, a relationship which we will call *interpretive resemblance* and which is easier to define.

In isolation, a proposition P (and, by extension, a representation with P as its content) has a number of analytic implications. However, propositions are entertained not in isolation but in a context of background assumptions. In a context {C}, a proposition P may have what we call *contextual implications*. A contextual implication of P in the context {C} is a proposition implied neither by {C} alone, nor by P alone, but by the union of {C} and P. We will say that two propositions P and Q (and, by extension, two representations with P and Q as their propositional content) interpretively resemble one another in a context {C} to the extent that they share their analytic and contextual implications in the context {C}.

We are thus defining interpretive resemblance as a context-dependent notion: two propositions P and Q may resemble one another closely in one context and less closely or not at all in another context. Let us briefly illustrate this with an artificially simple example. Consider:

- (4) It is winter.
- (5) It is freezing cold.
- (6) (a) If it is winter, then it is cold.  
(b) If it is cold, then we should stay at home.
- (7) (a) If it is winter, there are no flowers in the garden.  
(b) If it is freezing cold, we should heat the greenhouse.
- (8) It is cold.
- (9) We should stay at home.

By our definition, propositions (4) and (5) resemble one another more in context (6a–b) than in context (7a–b): in (6a–b) they share implication (8), which is contextually implied by (4) and analytically implied by (5), and implication (9), which is contextually implied by both; whereas in (7a–b), (4) and (5) share no implications at all. This seems to match our intuitions, insofar as intuitions are possible given the artificiality of the example.

Interpretive resemblance is a comparative notion with two extremes: no resemblance at all (i.e. no shared implications) at one end, and full propositional identity at the other. If two representations have the same propositional content, and hence share all their analytic implications, they also, of course, share all their contextual implications in every context. Let us say that when one representation is interpretively used to represent another, all of whose implications it shares, it is a *literal* interpretation of that other representation. On this account, literalness is just a limiting case of interpretive resemblance.

We began this section by suggesting that utterances can in principle represent something other than the states of affairs they describe. We now want to claim that they do. Every utterance used in verbal communication interpretively represents a thought entertained by the speaker—the very thought that the speaker wants to communicate. That much the hearer is entitled to expect; that much is necessary for verbal communication to be possible at all. However, the hearer is not invariably entitled to expect a literal interpretation of the speaker's thought, nor is such an interpretation always necessary for successful communication to take place. A less-than-literal interpretation of the speaker's thought may be good enough: may indeed be better on some occasions than a strictly literal one.

At this point, someone might raise the following objection. A rule of literal truthfulness at least explains how literal utterances are understood. Merely to assume that utterances interpret thoughts, without fixing the degree of resemblance required, may indeed put literal talk, loose talk, and metaphorical talk on a level, but it is a level on which all three are obscure. How is the hearer to assess the intended degree of resemblance? How is he to decide which of the implications of the utterance are shared with the speaker's thought, and which are not? This is, of course, a genuine problem. But before offering a solution, let us point out that it is one that already exists in the study of cognition in general, and of communication in particular.

A great many representations used by humans are representations in virtue of resemblance. Not all the properties of such a representation need, or generally even could, be shared by the original. For instance, when I draw you a diagram of how to get to my house, you do not infer that I intend you to travel across white paper, in two dimensions, past landmarks clearly labelled CHURCH and NEWSPAPER SHOP, a distance of 8 inches from door to door. You have to make some assumption about which properties of the representation carry over to the original. Or if I summarize in my own terms an article I have just read, you have to decide how close my summary is to the article, and in what respects.

So the problem of how intended resemblances are recognized is not an artifact of our approach. Any account of human communication will have to offer a solution to it. Since we believe we have a solution, we have no qualms about acknowledging the role of resemblance in communication.

### III

#### *Relevance*

Human information processing requires some mental effort and achieves some cognitive effect. Some effort of attention, memory and reasoning is required. Some effect is achieved in terms of alterations to the individual's beliefs: the addition of new beliefs, the cancellation of old beliefs, or merely a change in his degree of confidence in old beliefs. We may characterise a comparative notion of *relevance* in terms of effect and effort as follows:



(10) (a) Other things being equal, the greater the cognitive effect achieved by the processing of a given piece of information, the greater its relevance for the individual who processes it.

(b) Other things being equal, the greater the effort involved in the processing of a given piece of information, the smaller its relevance for the individual who processes it.

We claim that humans automatically aim at maximal relevance, i.e. maximal cognitive effect for minimal processing effort. This is the single general factor which determines the course of human information processing. It determines which information is attended to, which background assumptions are retrieved from memory and used as context, which inferences are drawn. Subjectively, of course, it seems that it is particular interests, transient or long-term, which guide our thoughts and determine the relevance of new information. We claim that interests are simply by-products of the general search for relevance: as a result of our cognitive history, some topics in our memory are richer in information and, either temporarily or permanently, more accessible than others, so that information relating to them is likely to produce greater effect for less effort, i.e. be more relevant as defined.

To communicate is, among other things, to claim someone's attention, and hence to demand some expenditure of effort. People will not pay attention unless they expect to obtain information that is rich enough in effects to be relevant to them. Hence, to communicate is to imply that the stimulus used (e.g. the utterance) is worth the audience's attention. Any utterance addressed to someone automatically conveys a presumption of its own relevance. This fact, we call the *principle of relevance*.

The principle of relevance differs from every other principle, maxim, convention or presumption proposed in modern pragmatics in that it is not something that people have to know, let alone learn, in order to communicate effectively; it is not something that they obey or might disobey: it is an exceptionless generalisation about human communicative behaviour. What people do have to know, and always do know when they recognize an utterance as addressed to them, is that the speaker intends that particular utterance to seem relevant enough to

them to be worth their attention. At this moment, you know that we intend this paper to seem relevant to you. In other words, what people have to recognize is not the principle of relevance in its general form, but the particular instantiations of it that they encounter.

Speakers may try hard or not at all to be relevant to their audience; they may succeed or fail; they still convey a presumption of relevance: that is, they convey that they have done what was necessary to produce an adequately relevant utterance.

Relevance, we said, is a matter of cognitive effect and processing effort. On the effect side, it is in the interest of hearers that speakers offer the most relevant information they have. However, speakers have their own legitimate aims, and as a result may choose to offer some other information which is less than maximally relevant. Even so, to be worth the hearer's attention, this information must yield at least adequate effects, and the speaker manifestly intends the hearer to assume that this is so. On the effort side, there may be different ways of conveying the same information, all equally easy for the speaker to produce, but requiring different amounts of processing effort from the hearer. Here, the speaker manifestly intends the hearer to assume that the formulation chosen is the one that is easiest to process. In other words, the presumption of relevance has two parts: a presumption of adequate effect on the one hand, and a presumption of minimally necessary effort on the other.

As is well known, linguistic structure grossly underdetermines the interpretation of an utterance: the linguistic meaning is generally ambiguous, it may be elliptical or vague, it contains referential expressions with undetermined referents, the intended illocutionary force is often not fully specified, and implicatures are not linguistically encoded at all. There are still other sources of underdetermination: one is that the strictly literal interpretation need not be the one intended, and, if we are right, is not even a preferred interpretation.

Various pragmatic theories appeal to complex sets of rules, maxims, or conventions to explain how this linguistic underdetermination is contextually overcome. We claim that the principle of relevance is enough on its own to explain how linguistic structure and background knowledge interact to determine verbal comprehension.

In a nutshell, for an utterance to be understood, it must have one and only one interpretation consistent with the fact that the speaker intended it to seem relevant to the hearer—adequately relevant on the effect side and maximally relevant on the effort side. We will say that in this case the interpretation is *consistent with the principle of relevance*, meaning consistent with the particular instantiation of the principle. The speaker's task is to make sure that the thought she intends to convey is consistent with the principle of relevance; otherwise, she runs the risk of not being properly understood. The hearer's task is to find the interpretation which is consistent with the principle of relevance; otherwise, he runs the risk of misunderstanding the utterance or not understanding it at all. In our book and several articles (see, for example, Sperber and Wilson 1981, 1982; Wilson and Sperber 1985, forthcoming), we have illustrated how this criterion of consistency with the principle of relevance works for various aspects of linguistic underdetermination. Here, we will show how it works to determine the intended resemblance between the utterance and the thought it is used to communicate.

#### IV

##### *Loose talk*

Loose talk is appropriate in the following, quite ordinary circumstances. The speaker wants to communicate to her hearer a certain set of propositions  $P_1 \dots P_n$ . They are all quite easily derivable as logical or contextual implications of a proposition  $Q$  whose truth she does not believe and does not want to guarantee. The best way of conveying this information may be to express the single proposition  $Q$ , as long as the hearer has some way of selecting those of its logical and contextual implications that the speaker intends to convey, and ignoring the others.

Our claim is that such a selection process is always at work: is part of the process by which every utterance is understood. Whenever a proposition is expressed, the hearer takes for granted that some subset of its logical and contextual implications are also logical or contextual implications of the thought being communicated, and aims to identify this subset. He assumes (or at least assumes that the speaker

assumed) that this subset will have enough cognitive effects to make the utterance worth his attention. He also assumes (or at least assumes that the speaker assumed) that there was no obvious way of achieving these effects with less processing effort. He aims for an interpretation consistent with these assumptions, i.e. consistent with the principle of relevance. When this criterion selects a single interpretation (or closely similar interpretations with no significant differences between them), communication succeeds.

Suppose Marie lives in Issy-les-Moulineaux, a block away from the city limits of Paris. At a party in London, she meets Peter. He asks her where she lives, and she answers:

(11) I live in Paris.

Marie's answer is literally false, but in ordinary circumstances it is not misleading. Peter will be able to infer from it a substantial amount of true or plausible information: that Marie spends most of her time in the Paris area, that she knows Paris, that she lives an urban life, that he might try to meet her on his next trip to Paris, and so on. It is such implications which make Marie's utterance relevant enough to be worth his attention, in a way Marie manifestly might have foreseen; moreover, there was no obviously more economical way of conveying these implications. Hence, Peter is entitled to assume that Marie intended him to interpret her utterance in this way, which is consistent with the principle of relevance.

Peter would be misled by Marie's answer only if he were to conclude from it that she lives within the city limits of Paris rather than in a suburb. However, it is clear that Marie had no reason to think that Peter would have to derive such a conclusion in order to establish the relevance of her utterance. Therefore her utterance does not warrant it.

Suppose, now, that Marie had answered instead:

(12) I live near Paris.

This time her answer is literally true, but it might well be misleading. The qualification '*near* Paris' demands some processing effort, which, given the presumption of relevance, should be offset by some cognitive effect. Peter might thus infer from this answer that Marie probably has to travel, say by

suburban transport, to get to Paris, that she lives a suburban life and so on, which is not the case. In other words, it is not just that Marie's first answer, 'I live in Paris', is effective enough to convey just what she wants; it may be more effective than the literally true second answer, 'I live near Paris'.

There will be cases where the subset of implications selected by the criterion of consistency with the principle of relevance will include the proposition actually expressed. Suppose Marie is asked where she lives, not at a party in London, but at an electoral meeting for a Paris local election. If she answers that she lives in Paris, the proposition expressed will itself be crucially relevant; hence the utterance will be understood literally, and Marie will have lied.

Our approach handles loose uses without abandoning truth-conditional semantics. If we are right, loose uses are non-literal uses in the sense described above: they are based on resemblance relations among representations, and involve interpretive rather than descriptive dimensions of language use. When a proposition or concept is loosely understood, it is not (or at least it need not be) that it is a vague concept or proposition; it is not that a guarantee of approximate truth is given to the proposition expressed: no guarantee of truth is given to this proposition at all. Instead, certain of its logical and contextual implications are taken to be accompanied by regular guarantees of truth, whereas others are simply ignored. Thus the truth-conditional relation between propositions and the states of affairs they represent remains unaltered: what varies is how closely the proposition expressed is taken to represent the speaker's thought.

We would like to suggest that much of the attraction of appeals to fuzzy concepts is that they seem to offer an account of data at least some of which might be better handled along the lines just described. The issue is whether we have well-defined classificatory concepts such that every object either does or does not fall under them, or whether our concepts are inherently fuzzy or open-ended, with no well-defined satisfaction conditions or clear boundaries between them. Does it make sense, for example, to ask whether a certain drinking vessel that seems to fall midway between being a cup and a mug is *really* a cup or a mug, or are these concepts fuzzy in such a way that either can

accommodate our vessel as a marginal case, low, so to speak, in cuppiness or mugginess?

Once the interpretive dimension of language use is taken into account, it is possible to suggest a rather different solution. In at least some cases, what is analysed as a literal use of a fuzzy concept might instead be analysed as a loose use of a classificatory concept, the looseness being motivated by the pursuit of relevance. In the above case, 'cup' and 'mug' may well have clear-cut boundaries and still be used loosely to refer to an object which falls outside these boundaries.<sup>1</sup>

This approach, incidentally, suggests a solution to the so-called baldness paradox. One is led into this paradox by agreeing, first, that a man with no hair is bald, next, that if a man with no hair is bald then a man with one hair is bald, and then, via the general principle that if a man with  $n$  hairs is bald then a man with  $n + 1$  hairs is bald, to the conclusion that a man with a full head of hair is bald. One way of avoiding this paradox is to treat *bald* as a classificatory concept with a necessary and sufficient condition: having no hair. Thus, to describe a man with one hair as bald is strictly speaking false, though of course perfectly appropriate as a loose use: virtually every conclusion that would be drawn from the fact that he was bald would apply to someone with only a single hair, and in most circumstances calling such person not bald, though literally true, would be grossly misleading. The same is true of somebody with very little hair. However, there comes a degree of looseness where not enough implications of calling someone bald are maintained. The point at which looseness becomes unacceptable varies with the context and is therefore not well defined in the abstract. On this approach, the claim that a man with one hair is bald is just as false as the claim that a man with a full head of hair is bald. What distinguishes them is not the fact that one is true and the other false, but the fact that one is an acceptable loose use because many of its logical and contextual implications are true, whereas the other is unacceptable since a hearer would be able to derive from it virtually no true descriptive information about the state of affairs it purports to represent.

<sup>1</sup>There are, however, independent reasons for thinking that not all utterances are fully truth-conditional.

## V

*Metaphor*

We want to claim that there is no discontinuity between loose uses and the most characteristic cases of poetic metaphor. In both instances the proposition expressed departs from full literalness. In both instances, however, the hearer must assume that the speaker is prepared to endorse some subset of the logical and contextual implications of the proposition expressed. In this way, information will be assertively conveyed by expressing a proposition which itself receives no guarantee of truth.

There is no discontinuity either between metaphor and a variety of other figures such as hyperbole, synecdoche or metonymy. Some of these provide intermediate cases between ordinary loose talk and typical metaphors. Consider an example of hyperbole. The speaker expresses the proposition in (13) and communicates a belief not in this proposition but in the weaker (14):

- (13) Bill is the nicest person there is.
- (14) Bill is a very nice person.

How can this be? Let us assume that by expressing (14) directly the speaker would not exhaust her thoughts about Bill: its contextual implications would fall short of what she wants to convey. Nor is there any obvious combination of adverbs and adjectives that would exactly express her thoughts. Perhaps they are too vague: there are a lot of aspects of Bill's niceness that she is not thinking about with equal clarity at the time, and to access these thoughts and make them more precise would involve more work than she is prepared to do. By expressing (13), she can be sure that all the propositions which make up the thought she does want to convey are among its contextual implications. Perhaps (13) has other contextual implications which she does not want to endorse. As long as she can rely on the hearer to ignore them, (13) will be a much more adequate representation of her thought than the weaker (14).

What exactly does (13) convey? The speaker is certainly guaranteeing the truth of (14), which is thus a strong implicature of (13). However, if this were all she had wanted to convey, she could have saved the hearer some processing effort

by expressing (14) directly. The greater effort imposed indicates that greater effect is intended. By uttering (13), the speaker thus encourages the hearer to look for a range of further contextual implications not shared by (14), and to assume that within this range there are some that she is prepared to endorse. He might conclude that the speaker finds Bill nicer than any of their common acquaintances; he might conclude that Bill has behaved in ways which show extraordinary niceness, and so on. Unlike (14), which is strongly implicated by (13), these further conclusions are only weakly implicated. That is, the hearer is encouraged to derive them, he can find some degree of confirmation for them in the utterance; however, this degree of confirmation may not be enough by itself and he must share the responsibility for deriving them. Thus (13) conveys a range of propositions, some, such as (14), very strongly and distinctly, others less so: a range of propositions which should closely resemble the complex thought that the speaker intended to share with the hearer. Conveying such a range of partly weak, partly strong implicatures is typical of the better-known figures of speech.

Metaphors vary in their degree of creativity. At one extreme are highly standardised examples such as (15):

(15) Jeremy is a lion.

Typically, such examples have one very strong implicature which constitutes the main point of the utterance: thus (15) implicates, in the context of stereotypical assumptions about lions, that Jeremy is brave. The fact that such metaphors are so regularly used with the same clearly defined implicature makes them relatively cheap to process, which in turn compensates for their relative poverty of content as compared to genuine creative metaphors. Nonetheless, they must suggest *some* further line of thought if their relative indirectness and its extra processing cost is to be justified. The speaker must be taken to have had in mind, however dimly, something more than Jeremy's bravery; and the hearer is encouraged to explore other contextual implications of (15), having to do, say, with the type of bravery Jeremy exhibits, or with his physical appearance. Thus even these highly standardised examples cannot be paraphrased without loss.



Our example (2) was that of a marginally more creative metaphor:

(2) *Mother to child*: You're a piglet.

While calling somebody a pig is quite standard, calling a child a piglet requires some extra processing effort, which should be offset by added effect. For instance, young animals are endearing, even when the adults of the species are not; so the child may feel encouraged to derive not only the obvious contextual implication that he is dirty, but also the further contextual implication that he is, nevertheless, endearing.

The most creative metaphors require of the hearer a greater effort in building an appropriate context, and deriving a wide range of implications. In general, the wider the range of potential implicatures and the greater the hearer's responsibility for constructing them, the more creative the metaphor. In the richest and most successful cases, the hearer can go beyond just exploring the immediate context and the background knowledge directly invoked, accessing a wider area of knowledge, entertaining ad hoc assumptions which may themselves be metaphorical, and getting more and more very weak implicatures, with suggestions for still further processing. The result is a quite complex picture, for which the hearer has to take a large share of the responsibility, but the discovery of which has been triggered by the speaker. The surprise or beauty of a successful creative metaphor lies in this extreme condensation, in the fact that a single expression which has itself been loosely used will determine a very wide range of acceptable weak implicatures.

For example, take Prospero's words to his daughter Miranda:

The fringed curtains of thine eyes advance  
And say what thou see'st yond.

(Shakespeare: *The Tempest* I ii)

Coleridge argues, against Pope and Arbuthnot, that these words should not be taken as equivalent in meaning to 'Look what is coming yonder'. They are uniquely appropriate to the characters and situation:

Prospero sees Ferdinand and wishes to point him out to his daughter not only with great but with scenic solemnity . . .

Something was to appear to Miranda on the sudden, and as unexpectedly as if the hearer of a drama were to be on the stage at the instant when the curtain is elevated . . . Turning from the sight of Ferdinand to his thoughtful daughter, his attention was first struck by the downcast appearance of her eyes and eyelids . . . (Coleridge: Notes on *The Tempest*)

Coleridge's comments are indeed illuminating, but they invite an objection and a question. The objection is that it is possible to appreciate Shakespeare's metaphor without understanding it exactly as Coleridge does. The question is how such an understanding is arrived at.

Our answer to the question also takes account of the objection. To understand Prospero's metaphor, the hearer must construct a context which will involve, on the one hand, his knowledge of the appearance of eyelids, and, on the other hand, his knowledge of curtains, and theatre curtains in particular. Merely retaining the implication that Prospero is telling Miranda to raise her eyelids—no doubt the strongest implicature—would result in an interpretation requiring too much effort for too little effect. A more creative hearer will invest a little more effort and get much more effect. This extra effort may consist in creating a metaphor of his own—for instance Coleridge's metaphor of the hearer of a drama being brought on stage—and adopting some of the joint implications of Prospero's metaphor and his. In such a process, the hearer is taking a large share of the responsibility for the conclusions he arrives at. As a result, different hearers with different background knowledge and different imaginations will follow somewhat different routes. However, they are all encouraged and guided by the text, and they all proceed by exploring its analytic and contextual implications as relevantly as they can.

In conclusion, let us consider how this approach to metaphor compares with the classical and romantic accounts. In many ways, we are on the romantic side. If we are right, metaphors are based on fundamental psychological mechanisms which are both natural and universal. They are in no sense departures from a norm or breaches of a rule or maxim of communication. They are simply creative and evocative exploitations of a basic

feature of all verbal communication: the fact that every utterance resembles, with a degree of closeness determined by considerations of relevance, a thought of the speaker's.

We also reject the classical claim that tropes in general, and metaphor in particular, have a purely decorative function. For us, as for the romantics, metaphor has a genuine cognitive content which, particularly with the more creative metaphors, is not paraphrasable without loss. This content we have proposed to analyse in terms of an indefinite array of weak implicatures whose recovery is triggered by the speaker, but whose content the hearer actively helps to determine.

Despite our general sympathy with the romantic view of metaphor, we differ sharply from the romantics on the nature of language and logic. We have tried to show that the existence of loose uses does not mean that language is irremediably fuzzy, and that the fact that language use is shot through with metaphor does not make metaphor an aspect of word and sentence meaning. Our aim has been to reconcile the view that looseness and metaphor belong to the most basic level of language use with a truth-conditional view of semantics.

Our main claim has been that hearers generally approach utterances without fixed expectations as to their literalness, looseness or metaphorical nature. They merely expect there to be an interpretive resemblance between the proposition expressed by the utterance and the thought that the speaker intends to convey. This expectation itself derives from, and is warranted by, a more basic expectation: an expectation of relevance. Such an expectation of relevance is automatically encouraged by any act of communication. This fact—the principle of relevance—is enough to explain how contextual information can be brought to bear on a linguistically underdetermined utterance, underdetermined in particular as regards its degree of literalness or looseness, and uniquely determine its interpretation.

#### REFERENCES

- Grice, H. P. 1975: Logic and conversation. In P. Cole and J. Morgan (eds) *Syntax and semantics 3: Speech acts*. New York: Academic Press, 41–58.
- Grice, H. P. 1978: Further notes on logic and conversation. In P. Cole (ed.) *Syntax and semantics 9: Pragmatics*. New York: Academic Press, 113–28.

- Sperber, D. and Wilson, D. 1981: Irony and the use-mention distinction. In P. Cole (ed.) *Syntax and semantics 9: Pragmatics*, New York: Academic Press: 295-318.
- Sperber, D. and Wilson, D. 1982: Mutual knowledge and relevance in theories of comprehension. In N. V. Smith (ed.) *Mutual knowledge*, London: Academic Press, 61-131.
- Sperber, D. and Wilson, D. 1986: *Relevance: Communication and cognition*. Oxford: Blackwell; Cambridge, Mass.: Harvard UP.
- Wilson, D. and Sperber, D. 1985: On choosing the context for utterance interpretation. In J. Allwood and E. Hjelmquist (eds): *Foregrounding Background*. Lund: Doxa, 51-64.
- Wilson, D. and Sperber, D. forthcoming: Inference and implicature. To appear in C. Travis (ed.) *Meaning and interpretation*. Oxford: Blackwell.