

5 A deflationary account of metaphors

Dan Sperber and Deirdre Wilson

5.1 Introduction

Are metaphors departures from a norm of literalness? According to classical rhetoric and most later theories, including Gricean pragmatics, they are. No, metaphors are wholly normal, say the Romantic critics of classical rhetoric and a variety of modern scholars ranging from hard-nosed cognitive scientists to postmodern critical theorists. On the metaphor-as-normal side, there is a broad contrast between those, like the cognitive linguists Lakoff, Talmy or Fauconnier, who see metaphor as pervasive in language because it is constitutive of human thought, and those, like the psycholinguists Glucksberg or Kintsch, or relevance theorists, who describe metaphor as emerging in the process of verbal communication.¹ While metaphor cannot be both wholly normal and a departure from normal language use, there might be distinct, though related, metaphorical phenomena at the level of thought, on the one hand, and verbal communication, on the other. This possibility is being explored (for instance) in the work of Raymond Gibbs.² In this chapter, we focus on the relevance-theoretic approach to linguistic metaphors.

Relevance theory's approach to metaphor is deflationary. Most rhetorical, literary and philosophical traditions emphasise both the importance and the distinctiveness of metaphor. We acknowledge its importance but dispute its distinctiveness. Metaphors are indeed ubiquitous in language use and contribute to what Barthes called 'le plaisir du texte'. Specific uses of metaphors by individual authors or in given literary genres are certainly worthy of study, and so is the very idea of metaphor as a culturally salient notion with a long and very rich history. Still, we see metaphors as simply a range of cases at one end of a continuum that includes literal, loose and hyperbolic interpretations. In our view, metaphorical interpretations are arrived at in exactly the same way as these other interpretations. There is no mechanism specific to metaphors, no interesting generalisation that applies only to them. In other terms, linguistic metaphors are not a natural kind, and 'metaphor' is not a theoretically important notion in the study of verbal communication. Relevance theory's account of metaphor is on the lean side, and is bound to disappoint those who feel that verbal metaphor deserves

a full-fledged theory of its own, or should be at the centre of a wider theory of language, or even of thought.

The widespread view that language use is governed by a norm of literalness (which is violated by metaphor and other figurative uses – hence their distinctiveness) follows straightforwardly from the even more widespread view that the function of language in communication is to allow the speaker to encode her meaning and the hearer to decode it. Loosening the grip of this ‘code model’ view of human communication is a necessary first step towards putting metaphor in a proper perspective.

5.2 The function of language in communication

A code is a systematic pairing of messages and signals. Encoding a message into a signal that a recipient can then decode is a very simple way to communicate very simple messages. Non-human animals do it all the time. Formally speaking, human languages are also codes: they are systems of sound–sense pairs generated by an underlying grammar. But although they are codes, human languages are vastly different from the codes of animal communication. First, and most obvious, they are incomparably richer. Languages not only contain a vast repertoire of expressive elements – the lexicon – with no counterpart in animal signalling systems, but these elements are combined by a syntax with unbounded generative capacities.

Human languages differ from animal codes in another respect that should be equally obvious but is hardly ever mentioned: they are grossly defective as codes. For communication to be achieved purely by coding and decoding, each signal in the code must unambiguously convey exactly the same content on all occasions. Ambiguity – where the same signal is paired with several messages – will stall the decoding process. True, there are cases even in animal communication where the precise message encoded by a given signal varies with the context. In the ‘bee dance’, for instance, the orientation of the bees’ communicative movements indicates the direction in which pollen is to be found, but this indication is relative to the position of the sun at the time. Limited context-sensitivity of this type can be handled by automatic code-like rules of disambiguation and accommodated in a coding–decoding system. However, the interpretation of the linguistic utterances that humans use to communicate is far too context-sensitive to be automatically achieved in purely code-like terms. The sentences of a natural language are typically multiply ambiguous; they contain referential expressions whose values cannot be assigned by decoding alone; the senses they ambiguously encode are often elliptical or incomplete; and there are still other ways in which the encoded meaning of a sentence falls short of determining what it may be used to communicate.

So although a language is formally a code, and human communication involves linguistic coding and decoding, there is a considerable gap between the semantic structure a sentence encodes and the meaning a speaker manages to convey by uttering that sentence in a given situation. In the case of metaphors and other tropes, this gap is often acknowledged as if it were an exception, and described in terms of a distinction between literal and figurative meaning.³ We claim that metaphors are not exceptional, and that the linguistic content of all utterances, even those that are literally understood, vastly underdetermines their interpretation.

When we say that human languages are defective as codes, we do not mean to imply that there is something wrong with them, or that we should want to improve on them (as some philosophers in the analytic tradition once proposed). On the contrary, we assume that human languages are exquisitely well suited to performing their function in communication. It is just that this function cannot be to encode speakers' intended meanings.

Humans communicate not only by using language but also by producing a variety of what we call 'ostensive stimuli': that is, actions (e.g. gestures or speech) or traces of actions (e.g. writings) that are manifestly intended to attract an addressee's attention and convey some content. Many of these ostensive stimuli do not belong to a code, and hence do not properly speaking encode anything. By using ostensive stimuli, humans are capable of communicating without language, and indeed without any other code. How can a stimulus convey a meaning that it doesn't encode? By providing evidence that the communicator intends to convey this meaning.

Suppose Mary is angry with Peter and doesn't want to talk to him. When he tries to engage her in conversation, she might

- (1) stare pointedly at the ceiling
- (2) open a newspaper and start reading it.

These actions do not draw on any established code. Still, what staring at the ceiling or opening a paper suggests to Peter is that Mary would rather do these things than talk to him at that time. Given that these actions are ostensive stimuli (i.e. are performed in order to attract his attention and convey some content to him), Peter understands Mary to *mean* that she doesn't want to talk to him. He interprets her in this way not because of some underlying code that systematically pairs stimuli of this type to a meaning of this type, but because her actions bring this interpretation to mind, and the best possible explanation of Mary's behaviour is to assume that this is just what it was intended to do. A stimulus can convey a meaning it does not encode by providing evidence that the communicator intends to convey this meaning. Here, the meaning is recovered not by decoding but by inference.

What is true of uncoded communicative stimuli is also true of coded stimuli used in human communication: they too convey their producer's intended meaning not by directly encoding it but by encoding some evidence of it. In the situation described above, Mary might

- (3) look angrily at Peter and clamp her mouth firmly shut
- (4) look angrily at Peter, put a finger to her lips, and whisper 'Shhh!'

In (3) and (4), Mary makes a gesture conventionally used to convey a request for silence, from which Peter can infer that she does not want to talk to him. Unlike the actions in (1) and (2), clamping one's mouth firmly shut or whispering 'Shhh' may be seen as encoding some meaning, but this encoded meaning is much vaguer than Mary's own meaning. For instance, the same gestures might be used in other situations to convey a request for secrecy. In the present situation, though, they are enough to indicate Mary's meaning.

In the same situation as before, Mary might also

- (5) say 'I am deaf and dumb'
- (6) say 'I won't talk to you'.

Obviously, the decoded linguistic content of Mary's utterance in (5) does not directly yield her meaning, but it provides a starting point for inferring her meaning that is not too different in effect from the gesture of clamping one's mouth shut, as in (3). In both cases, what is activated in Peter's mind is the idea of its being impossible to talk, an idea whose import is easy enough to work out in the situation.

What about Mary's utterance in (6)? Surely this, at least, encodes her exact meaning? In fact, it too falls some way short of doing so: the future tense does not indicate when Mary won't talk to Peter; the declarative sentence form does not indicate whether she is expressing a prediction, a warning, a threat or what. On another occasion, she might use the same sentence to promise Peter that she will talk to the whole group rather than just to him. Still, in the situation described, Peter can reconstruct Mary's full meaning by starting from the linguistic content of her utterance and specifying it further so as to reach a contextually plausible interpretation.

What these examples illustrate is the general point that, whether or not it involves the use of a language or some other code, human communication is inferential communication. The communicator provides some evidence of her meaning and the addressee infers this meaning on the basis of this evidence and the context. The evidence may or may not be coded, and if it is coded, it may or may not be linguistic, but in each case it provides input to an inferential process whose goal is to understand the communicator's meaning. Which raises the following question: What is the point of using a language at all, if the kind of

thing it can be used to achieve can also be achieved without it? The point is that a language provides an unbounded repertoire of evidence of the speaker's meaning, evidence that can be as nuanced, as complex, as richly structured as the speaker likes. Non-verbal kinds of evidence are much more limited. With language (and only with language) people can communicate about anything they can think about, whether they can point to it or not, imitate it or not, and they can do this with endless refinement. The fact that the interpretations of utterances are not encoded but merely evidenced by their linguistic meaning does not detract from the richness of linguistic communication, but, on the contrary, enhances it: every single sentence can give rise to an open array of interpretations which go well beyond the encoded senses. Some of the best illustrations of this are, of course, creative metaphors.

5.3 How relevance guides inferential comprehension

What we have sketched so far is a view of verbal communication suggested by the work of the philosopher Paul Grice, but more radical than his. Grice characterised a speaker's meaning as an overt intention to cause a certain cognitive effect in an audience via their recognition of one's intention to cause this effect (Grice 1989: chapters 5, 6, 14, 18). A speaker's meaning, so understood, is an intention, a mental state. The mental states of others cannot be simply perceived or decoded, but must be inferred from their behaviour, together with background information. What is special about a speaker's meaning as compared with other mental states (which people usually keep to themselves) is that speakers intend their audience to discover their meaning, and provide evidence to that effect, in the form of communicative behaviour. This raises the possibility that there might be an inferential procedure uniquely adapted to comprehension.

Grice tended to take for granted – and Searle explicitly argued – that when someone uses language to communicate, she is presumed to express her meaning literally. It can then be assumed by default that the literal linguistic meaning of the utterance is *her* meaning, or at least the explicit part of her meaning (Grice's *what is said*), with only the implicit part (Grice's *implicatures*) left to be inferred. This amounts, in practice, to saying that part of the speaker's meaning is decoded and part is inferred. Metaphors and other tropes, where the linguistic meaning of the utterance is not even part of the speaker's meaning, are exceptional in this respect: Grice suggested that in metaphor, the speaker is not really saying what she appears to be saying, but merely 'makes as if to say' it, so that in this case, the speaker's meaning must be wholly inferred. We claim, by contrast, that verbal comprehension involves no presumption of literalness and no default interpretation, and that metaphors are in no way exceptional. All human intentional communication works in the way outlined above: the communicator produces a piece of evidence of her meaning – the ostensive stimulus – and

the addressee infers her meaning from this piece of evidence and the context. Linguistic utterances are just one type of ostensive stimulus. Verbal communication is always context-sensitive and inferential.

How exactly does inferential comprehension work? Relevance theory draws on a precise characterisation of relevance and its role in human cognition to put forward a testable account of the comprehension mechanism, an account in which expectations of relevance play a crucial role.

We analyse *relevance* not just as a property of utterances or other ostensive stimuli, but as a property that any input to a cognitive process might possess: sights, sounds, utterances, thoughts, memories, suppositions may all be relevant to an individual at a given time. When is an input relevant? When processing it in the context of previously available information yields new cognitive effects. The input may answer a question the individual had in mind, it may raise or settle a doubt, suggest a hypothesis or a course of action, confirm or disconfirm a suspicion, correct a mistake. All these cognitive effects involve a fruitful interaction between the input and the context in which it is processed. However, the interaction may be more or less fruitful; inputs may be more or less relevant.

What makes one input more relevant than another? Suppose you are a caterer making lunch for a group of ten people, and all you need to know is how many will want the vegetarian menu. Then the information that three of them are vegetarian would be more relevant to you than the information that three of them are Buddhists (from which it follows that they are probably, though not definitely, vegetarian). In general, it is more informative to discover that someone is a Buddhist than to discover that he is a vegetarian, but if the context is such that only his food preferences are consequential, then the less informative input is more relevant. The greater the cognitive effects produced by processing an input, the greater its relevance (to the person processing it, at the time).

However, cognitive effects are only one of two factors that affect the relevance of an input. The other is the processing effort involved in achieving these effects. Some effort of perception, memory or inference is required to represent the input, access contextual information, and derive cognitive effects. In the situation described above, suppose that the choice is between a straightforward statement that three of the guests are vegetarian and a brochure with a short biography of all ten guests, mentioning *inter alia* whether or not they are vegetarian. In this case, the brochure would be less relevant than the straightforward statement: although both would contain all the information required, extracting this information from the brochure would involve more effort for the same effect, hence less relevance. In a nutshell:

Degrees of relevance:

- (a) The greater the *cognitive effects* achieved by processing an input, the greater its relevance.
- (b) The smaller the *processing effort* required to achieve these effects, the greater the relevance.

At every moment in their waking lives, humans have a huge variety of inputs competing for their attention: things and events they perceive, previous thoughts that have not been fully digested, pending goals, and so on. For contexts to use in processing these inputs, they have a vast mental encyclopaedia of accumulated knowledge on which to draw. At any given moment, most of these inputs are not worth processing, and, for any given input, most of this background information is not worth activating: the resulting process would yield too few cognitive effects to be worth the effort. Cognitive efficiency is very much a matter of selecting the most relevant inputs available at each point, and processing them in the context of background information that will most enhance their relevance. In fact, if there were not a strong tendency to select maximally relevant inputs, cognition would be an extremely wasteful activity. We assume that, among the many selective pressures that have driven the evolution of human cognitive capacities, there has been a constant pressure on the cognitive system as a whole, on its component parts, and on their articulation, towards an efficient use of brain resources. We therefore put forward the following claim:

Cognitive Principle of Relevance:

Human cognition tends to be geared to the maximisation of relevance.

We are not claiming that humans always succeed in maximising relevance, but only that they have a sufficient tendency to do so to make their massive investment in cognition evolutionarily worthwhile. More specifically, we are claiming that human perceptual mechanisms tend to pick out potentially relevant stimuli, human retrieval mechanisms tend to activate potentially relevant background assumptions, and human inferential mechanisms tend to process them in the most productive way, so that, overall, attention tends to go to the inputs with the greatest expected relevance. These claims have a variety of experimentally testable consequences (see Van der Henst and Sperber 2004). Here we are only concerned with the consequences of the cognitive principle of relevance for human communication.

Given the indefinite variety of possible objects of attention and courses of thought, it would be impossible for one person to predict what others will attend to, and what the resulting thoughts will be, if their attention and thought processes were not guided by considerations of relevance. The tendency to maximise relevance is crucial to making human mental processes relatively interpretable and predictable. Given this tendency, it is possible not only to interpret and predict the mental processes of others, but also to manipulate them by producing a stimulus which will predictably attract their attention and be interpreted in foreseeable ways. Jill knows it is relevant to Peter that all his guests should be happy, so she leaves her empty glass in his line of sight, anticipating that he will pay attention and conclude that she would like another drink. This is not yet a case of inferential communication, because, although Jill intends Peter to come to this conclusion, she provides evidence only that she is thirsty, and not that she intends to inform Peter that she is thirsty. If she had instead established eye contact with him and waved her empty glass, or said to him 'My glass is empty', then the stimulus would be ostensive, and her behaviour would be properly communicative.

Use of an ostensive stimulus as opposed to a regular non-ostensive one provides the addressee with information not only about some state of affairs (e.g. the fact

that Jill would like another drink) but also about the communicator's intention to convey this information, and to do so overtly. By producing an ostensive stimulus, the communicator openly requests the addressee's attention. Since attention tends to go to the most relevant inputs available, the communicator implicitly conveys that her message is such an input. The central claim of relevance-theoretic pragmatics is that use of an ostensive stimulus raises expectations of relevance not raised by other inputs, and that these expectations guide the comprehension process. More specifically, we claim:

Communicative Principle of Relevance

Every act of inferential communication conveys a presumption of its own optimal relevance.

The presumption of optimal relevance mentioned in the communicative principle has a precise content. The utterance (or other communicative act) is presumed to be relevant enough to be worth processing, from which it follows that it must be more relevant than other inputs competing for the addressee's attention at the time. In some conditions, it can be presumed to be even more relevant than that. Communicator and addressee have at least one common goal: for communication to succeed, that is, for the addressee to understand what the communicator meant. The more relevant the utterance, and in particular the less processing effort it requires, the more likely the addressee is to understand it successfully. The communicator can therefore be expected, within the limits of her expressive abilities, and without going against her own goals (and in particular the goal she is pursuing in communicating), to have aimed at maximal relevance. So when we say that every act of inferential communication conveys a presumption of its own optimal relevance, we mean something quite precise: as much relevance as is compatible with the communicator's abilities and preferences, and, in any case, enough relevance to be worth processing.

The communicative principle of relevance suggests both a path for the addressee to follow in constructing the interpretation of an utterance, and a stopping point. Since effort is one of the two factors affecting relevance, the appropriate path to follow is one of least effort. The stopping point is the point at which the current interpretation (i.e. what the speaker is taken to have conveyed, either explicitly or implicitly) satisfies the expectations of relevance raised by the utterance itself. From the speaker's point of view, the easiest way to increase the relevance of her communication, and hence the chances of being properly understood, is to express herself (within the limits of her abilities and preferences) so that the interpretation she intends to convey is the first interpretation found by an addressee following the path of least effort that meets the expectations of relevance she herself has raised.

To illustrate, consider the following exchange:

- (7) PETER: For Billy's birthday party, it would be nice to have some kind of show.
 MARY: Archie is a magician. Let's ask him.

Suppose that 'magician' is ambiguous for Peter, with two senses: (a) *someone with supernatural powers who performs magic*, and (b) *someone who does*

magic tricks to amuse an audience. In the context of a discussion about a show for a child's birthday party, the second sense is likely to be activated first, and the information (or reminder) that their friend Archie is a magician in this sense is likely to satisfy Peter's expectations of relevance by implying that he might perform at Billy's birthday party. In presuming that her utterance would be relevant to Peter, Mary must have expected him to derive this implication, which can therefore be seen as an implicit part of her meaning, i.e. an implicature. The disambiguation of 'magician' as someone who does magic tricks dovetails with this implicature, and the two confirm one another by jointly yielding an interpretation that is relevant in the expected way.

The linguistic meaning of the sentence 'Let's ask him' is very schematic and gappy, leaving the second part of Mary's utterance wide open to an indefinite range of interpretations. 'Him' may refer to Archie, to Billy, or to someone else. 'Ask' may be understood as asking for advice, help, an opinion, a favour, and so on. Thus, the whole sentence might be used to mean *Let's ask Billy whether he would like to have Archie perform magic tricks at his birthday party.* This interpretation would make sense in the situation, and would be quite compatible with Grice's maxims of conversation, or with standard theories of discourse coherence. Still, in a context where the first part of the utterance ('Archie is a magician') implicates that Archie could perform magic tricks at Billy's party, the first interpretation found by following a path of least effort will be that Peter and Mary should ask Archie to perform. Since this would satisfy Peter's expectations of relevance, he should accept it as the intended interpretation, without looking any further for alternative interpretations that might also be relevant. (None of these other potential interpretations could be optimally relevant, because extra processing effort would be required to retrieve them. They are therefore not worth considering unless there is some reason to think that Mary has failed to express herself in an optimally relevant way.)

In this example, Mary is speaking literally (which shows how far even the interpretation of an utterance that is literally understood can go beyond its linguistic meaning and is not just a simple matter of decoding). Our claim is that the very same procedure that yields a literal interpretation in this case would yield a non-literal interpretation in others.

5.4 Meaning construction

The decoded senses of a word or other linguistic expression in an utterance provide a point of departure for an inferential process of meaning construction. The meaning constructed may be narrower than the decoded meaning, as in (8) or (9):

- (8) I have a temperature.

- (9) PETER: Does Gérard like eating?
MARY: He's French!

In (8), 'temperature' would be understood as meaning a temperature above normal.⁴ What the speaker is communicating would be false if her temperature were a regular 37°C/98.6°F. In (9), what Mary means is not just that Gérard is a French national, but that he is what she regards as a prototypical Frenchman, and therefore someone who likes eating.

On other occasions, the meaning constructed may be broader than the decoded meaning, as in (10)–(14):

- (10) Holland is flat.
(11) The stones form a circle.
(12) (*On a picnic, pointing to a flattish rock*): That's a table!
(13) (*Handing someone a tissue*): Here's a Kleenex.
(14) (*Handing someone a paper napkin*): Here's a Kleenex.

The uses of 'flat' in (10) and 'circle' in (11) are cases of approximation. Approximation is a variety of loose use or broadening in which a word with a relatively strict sense is extended to a penumbra of items (what Lasersohn 1999 calls a 'pragmatic halo') that strictly speaking fall outside its linguistically specified denotation. The uses of 'table' in (12) and 'Kleenex' in (13) and (14) are cases of category extension. Category extension, another variety of loose use or broadening, involves applying a word with a relatively precise sense to a range of items that clearly fall outside its linguistically specified denotation, but that share some contextually relevant properties with items inside the denotation. Thus, the flat rock referred to in (12) is definitely not a table, but has properties which make it a good substitute for a table on that occasion. The tissue referred to in (13) is not a Kleenex, but will do just as well. The paper napkin referred to in (14) is not even a tissue, but is the closest available thing to a tissue, and will do almost as well.

With narrowing, literalness is in some sense preserved: a high temperature is literally a temperature, and a Frenchman who likes eating is literally a Frenchman. With broadening, literalness is not preserved: Holland is not literally flat, the stones do not literally form a circle, the flattish rock is not literally a table, and neither the tissue nor the paper napkin is literally a Kleenex. However, narrowing and broadening are not two functionally distinct types of language use. They both involve the same process of meaning construction, which happens in some cases to lead to a narrowing of the encoded concept, and in other cases to a broadening.

How are these narrowed or broadened lexical meanings arrived at? By following the relevance-guided comprehension procedure outlined above. With (8) ('I have a temperature'), a literal interpretation based on the decoded meaning of

'temperature' would be an irrelevant truism, since anyone (or indeed anything) has a temperature, just as it has a mass or a location. In fact, there is no reason to think that the hearer constructs and entertains such a truism. Rather, what happens is that the concept TEMPERATURE is activated in the hearer's mind and points him towards a relevant interpretation. This concept is (or has) a parameter that can take a range of values, some of which would be relevant in the circumstances (by implying, for instance, that the speaker is ill and unable to work). In the process of arriving at a relevant overall interpretation of the utterance, the decoded concept TEMPERATURE provides a starting point for constructing a narrowed *ad hoc* concept TEMPERATURE* which ranges only over contextually relevant temperatures: that is, temperatures which depart from the human norm in a way that is easily brought to mind, with implications that are worth the hearer's processing effort.

Similarly, activation of the lexicalised concept FLAT in (10) ('Holland is flat') gives access to a range of implications that would follow from Holland's being strictly flat: that it is a good place for easy cycling or not a good place for mountaineering, for instance. These implications hold (to different degrees for different implications) even if Holland is only approximately flat. In a context where (10) is relevant, some of these implications will be immediately obvious to the hearer, and will fulfil his expectations of relevance. The resulting overall interpretation (including the presumption of relevance and the implications that make the utterance relevant) will be internally consistent on the assumption that 'flat' in (10) indicates the speaker's intention to convey that Holland is FLAT*, where the *ad hoc* concept FLAT* represents an approximation to flatness which is close enough to yield the implications that make the whole utterance contextually relevant (for a detailed discussion of this and related examples, see Wilson and Sperber 2002).

In these two examples, the words 'temperature' or 'flat' are used in an utterance to evoke (or, more technically, to activate to some degree) potential implications of the encoded concepts TEMPERATURE or FLAT. More generally, we claim that ideas evoked in comprehension stand in inferential relationships to the concepts that evoke them,⁵ and are not mere associations based on past co-occurrence, with no inferential status. That is, the ideas evoked by the presence of a word in an utterance are likely to be true of items in the linguistically specified denotation of the word, or, equivalently, of items in the extension of the concept encoded by the word. In the case of narrowing, the implications hold across only part of the extension of the encoded concept (e.g. only some temperatures imply illness). In the case of broadening, the implications hold not only of items in the extension of the encoded concept but also of contextually salient items which fall outside the extension, but which share with items inside the extension properties that determine these implications (e.g. cycling is easy not only in flat, but also in flattish terrains).

Some of the implications evoked by the presence of a word are simultaneously evoked by the context. In (13) and (14) ('Here's a Kleenex,' said of a tissue or a paper napkin), implications of the form ... *can be used to blow one's nose* are activated in the hearer's mind not only by the word 'Kleenex' but by the fact that he has just been sneezing. Implications activated by both the utterance and the context are the first to come to mind, and are tentatively added to the interpretation until the hearer's expectations of relevance are satisfied. At that point, the explicit content of the utterance (in the case of an assertion, the propositions whose truth the speaker is committing herself to) is retroactively determined by mutually adjusting the implicit and explicit components of the interpretation. The explicit content of an utterance must be such that it contextually implies the implicit content. More technically, and in relevance-theoretic terms, the explicatures of an utterance must be such that, together with the implicit premises of the utterance, they warrant the derivation of its implicit conclusions (where both implicit premises and implicit conclusions are kinds of implicature). (On the mutual adjustment process, see Sperber and Wilson 1998a, 2005; Carston 2002a; Wilson and Sperber 2002, 2004; Wilson and Carston 2007.)

In the case of (8) ('I have a temperature'), the result of the mutual adjustment process is a contextual construal of 'temperature' as TEMPERATURE*, which is narrower than the lexicalised concept TEMPERATURE. In the case of (10) ('Holland is flat'), the result is a contextual construal of 'flat' as FLAT*, which is broader than the lexicalised concept FLAT. Narrowings and broadenings of meaning are thus arrived at by exactly the same procedure of online concept construction, and for the same reasons. In fact, as noted by Carston (1997), they may be combined in a single construal. Suppose that Mary in (9) says of Gérard, 'He's French!', intending to implicate that he likes eating, when, in fact, she knows that Gérard happens to be a citizen of Monaco. She would then be using neither the concept FRENCH, which denotes French nationals and is encoded (let us assume) by the word 'French', nor an appropriate narrowing, FRENCH*, but a concept FRENCH** which is narrower in some respects and broader in others, denoting people who fit some prototype of a French person without French nationality being either a sufficient condition or an absolutely necessary one for inclusion in its extension.

Strictly literal interpretations – those that involve neither narrowing nor broadening of the lexicalised concept – are arrived at by exactly the same process of mutually adjusting explicit content with implicit content. A literal interpretation results when the implications that make the utterance relevant in the expected way depend on the presence in the explicit content of the lexicalised concept itself (rather than some broadened or narrowed version).⁶ Literal interpretations are not default interpretations: they are not the first to be considered, and they are not necessarily easier to construct than non-literal ones. In fact, some literal interpretations are fairly hard to get, as in (15):

- (15) If Holland were flat, water would flow from the borders towards the centre.

In describing a stretch of land as ‘flat’, we broaden the concept by ignoring not only the various types of unevenness present in any terrain, but also the curvature of the earth.⁷ This second departure from the literal meaning of ‘flat’ is not so easily corrected.

There is a continuum of cases between approximations such as (10) and (11) and hyperboles. In fact, the same utterance can be properly understood hyperbolically, loosely, or literally, depending on the facts of the matter, with no sharp dividing line between the different interpretations. Consider (16):

- (16) MARY TO PETER: The soup is boiling.

If Peter is too far away to observe the state of the soup directly, how is he to select one of these possible interpretations? Via considerations of relevance. Suppose he is working upstairs; when he smells the soup Mary is making and says he is coming down to taste it, Mary answers as in (16). Then her utterance would be relevant as a warning not to bother: ‘boiling’ would function as a hyperbole, conveying *too hot to taste*. Or suppose instead that Peter is making the soup but has left the room, and Mary knows that the soup should not be allowed to boil at this stage. Then her utterance would be relevant enough if the soup were *almost* boiling: a loose, approximate use rather than a hyperbole. Suppose, finally, that Peter is making the soup but has left the room, and Mary knows that he wanted to skim it once it was properly boiling. Then in order to be relevant enough, her utterance would have to be interpreted literally.

5.5 The literal–loose–metaphorical continuum

There is a continuum of cases between limited category extensions such as (12)–(14) above and more creative ones such as (17) and (18):

- (17) Žižek is another Derrida.
 (18) For luggage, pink is the new black. (*New York Times*, 4 September 2005)

In (17) ‘Derrida’ is used as a common noun to denote a category of flamboyant and obscure philosophers à la Derrida. In (18), ‘black’ is used to denote a category of fashionable colours. In both cases, a category is extended to include items that share with its members some properties which may or may not be essential, but are at least salient. These examples of category extension, unlike the use of ‘Kleenex’ to refer to any tissue, are not analysable as mere loose uses. The claim in (17) is not that the differences between Žižek and Derrida are inconsequential, but that Žižek belongs to a broader category of which Derrida is the most salient member. The claim in (18) is not that pink is pretty much the same as black, but that it occupies, in the category of colours for luggage, the

place previously occupied by black. Still, (17) and (18) are interpreted by the usual process: the presence of the words ‘Derrida’ or ‘black’ helps to activate implications about Žižek, on the one hand, and the colour pink, on the other, that make the utterance relevant in the expected way. By mutual adjustment of explicit content and implicatures, the explicit content is construed as containing an *ad hoc* concept (DERRIDA* or BLACK*) that contextually carries these implications.

There is a continuum of cases between hyperbole and metaphor. It might seem at first blush that hyperbole involves only a quantitative difference between the concept encoded and the concept contextually constructed, as in (19) below, while metaphor also involves a qualitative difference, as in (20):⁸

(19) Joan is the kindest person on earth.

(20) Joan is an angel.

However, the quantitative–qualitative distinction is not sharp. For instance, (21) and (22) would generally be classified as hyperboles rather than metaphors, although there is both a quantitative and a qualitative difference between something that is credible and something that is not, or between a saint and an ordinary kind person:

(21) Joan is incredibly kind.

(22) Joan is a saint.

In any case, whether they are classified as hyperboles or metaphors, (21) and (22) would be interpreted in the same way: the encoded concept helps to activate contextual implications that make the utterance relevant as expected, and the concept conveyed by the hyperbole/metaphor is one of an outstanding type of kindness characterised by these implications.

There is also a continuum of cases between category extension and metaphor. It might be argued that category extension involves the projection of defining, or at least characteristic, properties of the encoded concept onto a broader category, as in (12)–(14) and (17)–(18) above, whereas the type of broadening involved in metaphor is based on relatively peripheral, or at least contingent, properties, as in (23) or (24):

(23) ‘Man is but a reed, the weakest in nature.’ (Blaise Pascal)

(24) My mind is cloudy.

Weakness is not a defining property of reeds (and it is only a property relative to some arbitrary comparison class); similarly, the difficulty of discerning parts is not a defining property of clouds.

However, some metaphors are based on fairly central properties of the lexicalised category. For instance, when the term for an animal body part is extended to a human body part, as in (25), the result would generally be classified as a metaphor:

(25) Henry was proud of his mane.

A category may undergo successive broadenings, with more peripheral extensions necessarily losing some of the most central features of the lexicalised category. Thus, compare (17) ('Žižek is another Derrida') with (26) and (27):

(26) Rebecca Horn is the Derrida of contemporary art.

(27) Ferran Adria is more Derrida than Danko. (attested: www.egullet.org/tdg.cgi?pg=ARTICLE-tabledancingadria – Adria is the world famous chef of El Bulli, Danko is a famous San Francisco chef)

In each case, a different concept (DERRIDA*, DERRIDA**, DERRIDA***) is constructed, each marginally further away from the original concept (if we accept that there are concepts of individuals) or representation of Jacques Derrida.

Central and peripheral properties may combine, as in (28), a comment on a clip of George W. Bush allegedly wiping his glasses on an unsuspecting woman's shirt during an appearance on Jay Leno's TV show:

(28) We're all human Kleenex to him (attested: www.ifiipflop.com/2004/10/metaphor-george-bush-uses-woman-as.html).

Here, the woman is implicitly described as a Kleenex, since she (or at least her clothes) can be used as one, and this carries the suggestion that Bush sees people as disposable artefacts of little value.

Most hyperboles involve only broadening of the encoded concept, with no narrowing. In (19), for instance, 'the kindest person on earth' (despite its singular form) is broadened to cover all very kind people, including Joan. By contrast, most metaphors involve both narrowing and broadening, and so cannot be seen simply as cases of category extension. In the metaphorical (20), 'angel' is interpreted as ANGEL*, which is narrowed, on the one hand, to cover only prototypical kind, caring angels (excluding avenging angels, angels of wrath or fallen angels) and broadened, on the other, to cover all very kind, caring people. However, this combination of narrowing and broadening is not a defining feature of metaphor. In the metaphorical (28), for instance, 'Kleenex' is broadened to something like the category of DISPOSABLE ITEMS, and this includes not only prototypical Kleenex but all Kleenex.

5.6 Inferential steps

We see this continuity of cases, and the absence of any criterion for distinguishing literal, loose, and metaphorical utterances, as evidence not just that there is

Table 5.1. *Interpretation of Mary's utterance, 'Archie is a magician'*

(a) Mary has said to Peter 'Archie is a magician'.	<i>Decoding of Mary's utterance.</i>
(b) Mary's utterance is optimally relevant to Peter.	<i>Expectation raised by the recognition of Mary's utterance as a communicative act.</i>
(c) Mary's utterance will achieve relevance by addressing Peter's suggestion that they have a show for Billy's birthday party.	<i>Expectation raised by (b), given that Mary is responding to Peter's suggestion.</i>
(d) Magicians (in one lexicalised sense of the term, MAGICIAN ₂) put on magic shows that children enjoy.	<i>Assumption activated both by use of the word 'magician' and by Peter's wish to have a show for Billy's birthday party. Tentatively accepted as an implicit premise of Mary's utterance.</i>
(e) Archie could put on a magic show for Billy's birthday party.	<i>Implicit conclusion derivable from (d), together with an appropriate interpretation of Mary's utterance, which would make her utterance relevant-as-expected. Tentatively accepted as an implicit conclusion of the utterance.</i>
(f) Archie is a MAGICIAN ₂ .	<i>Interpretation of the explicit content of Mary's utterance as decoded in (a) which, together with (d), would imply (e). Interpretation accepted as Mary's explicit meaning.</i>
(g) Archie is a MAGICIAN ₂ who could put on a magic show for Billy's birthday party that the children would enjoy.	<i>First overall interpretation of Mary's utterance (explicit content plus implicatures) to occur to Peter which would satisfy the expectation of relevance in (b). Accepted as Mary's meaning.</i>

some degree of fuzziness or overlap among distinct categories, but that there are no genuinely distinct categories, at least from a descriptive, psycholinguistic or pragmatic point of view.⁹ Even more important than the lack of clear boundaries is the fact that the same inferential procedure is used in interpreting all these different types of utterance. Let us look in more detail at how this procedure applies to the interpretation of two examples, one at the literal end of the continuum, and the other at the metaphorical end.

At the literal end, we return to example (7):

- (7) PETER: For Billy's birthday party, it would be nice to have some kind of show.
 MARY: Archie is a magician. Let's ask him.

Table 5.1 shows the inferential steps that Peter goes through in interpreting the first part of Mary's utterance ('Archie is a magician'), with Peter's interpretive hypotheses on the left, and his basis for arriving at them on the right.

At the metaphorical end of the continuum, consider (29):

- (29) PETER: I've had this bad back for a while now, but nobody has been able to help.
 MARY: My chiropractor is a magician. You should go and see her.

Table 5.2. *Interpretation of Mary's utterance, 'My chiropractor is a magician'*

(a) Mary has said to Peter 'My chiropractor is a magician'.	<i>Decoding of Mary's utterance.</i>
(b) Mary's utterance is optimally relevant to Peter.	<i>Expectation raised by the recognition of Mary's utterance as a communicative act.</i>
(c) Mary's utterance will achieve relevance by addressing Peter's expressed concern about his back pain.	<i>Expectation raised by (b), given that Mary is responding to Peter's complaint.</i>
(d) Chiropractors are in the business of healing back pain.	<i>Assumption activated both by use of the word 'chiropractor' and by Peter's worry about his back pain. Tentatively accepted as an implicit premise of Mary's utterance.</i>
(e) Magicians (in one lexicalised sense of the term, MAGICIAN ₁) can achieve extraordinary things.	<i>Assumption activated both by the use of the word 'magician' and by Peter's worry that no ordinary treatments work for him. Tentatively accepted as an implicit premise of Mary's utterance.</i>
(f) Mary's chiropractor, being in the business of healing back pain and able to achieve extraordinary things, would be able to help Peter better than others.	<i>Implicit conclusion derivable from (d) and (e), together with an appropriate interpretation of Mary's utterance, which would make her utterance relevant-as-expected. Tentatively accepted as an implicit conclusion of the utterance.</i>
(g) Mary's chiropractor is a MAGICIAN* (where MAGICIAN* is a meaning suggested by the use of the word 'magician' in the sense of MAGICIAN ₁ and enabling the derivation of (e)).	<i>Interpretation of the explicit content of Mary's utterance as decoded in (a) which, together with (d) and (e), would imply (f). Interpretation accepted as Mary's explicit meaning.</i>
(h) Mary's chiropractor is a MAGICIAN*, who would be able to help Peter better than others by achieving extraordinary things.	<i>First overall interpretation of Mary's utterance (explicit content plus implicatures) to occur to Peter which would satisfy the expectation of relevance in (b). Accepted as Mary's meaning.</i>

Table 5.2 shows, again in simplified form, the inferential steps that Peter goes through in interpreting the first part of Mary's utterance ('My chiropractor is a magician').

In both cases, of course, interpretation is carried out 'online', and starts while the utterance is still in progress. We assume, then, that interpretive hypotheses about explicit content and implicatures are developed partly in parallel rather than in sequence, and stabilise when they are mutually adjusted so as to jointly confirm the hearer's expectations of relevance. And we are not suggesting that the hearer consciously goes through just the steps shown in the tables, with exactly those premises and conclusions. We are not making claims about exact sequences, consciousness, or the representational format of thought. We are making claims about the factors that cause hearers to converge on an interpretation

which – when communication is successful – coincides with the one intended by the speaker.

Although ‘magician’ is interpreted literally in (7) and metaphorically in (29), the same kind of process is involved in both cases. With (7), the fact that one of the lexicalised senses of ‘magician’ is *MAGICIAN₂*, *someone who performs magic tricks to amuse an audience*, makes it particularly easy to access implications associated to this interpretation. Since these implications end up satisfying the hearer’s expectations of relevance and are carried only by this precise meaning, one of the lexicalised senses of ‘magician’ is selected by the comprehension process as the contextually indicated meaning. With (29), ‘magician’ provides easy access to the information that if someone is a magician, they have extraordinary capacities, and this is enough to ground an optimally relevant overall interpretation. The concept used in this interpretation is substantially broader than *MAGICIAN₁*, so in this case, as a rhetorician would say, ‘magician’ is a metaphor. However, the hearer pays no more attention to the fact that ‘magician’ is used metaphorically in (29) than he does to the fact that it is used literally in (7).

For that matter, some people may have only a single encoded sense for ‘magician’: *someone with supernatural powers who performs magic*. They would still have no difficulty arriving at an appropriate interpretation of (7) by extending the category of ‘real’ magicians to include make-believe ones. For other people, the metaphorical sense may have become lexicalised, so that ‘magician’ now has the additional encoded sense *someone who achieves extraordinary things*. They would obviously have no trouble arriving at an appropriate interpretation of (29). Mary did not intend her utterance to be understood literally in (7) and metaphorically in (29); her communicative intentions – like those of all speakers – are about content and propositional attitude, not rhetorical classification.

Relevance theory’s resolutely inferential approach to comprehension suggests a solution to the ‘emergent property’ issue raised in recent work on metaphor.¹⁰ Consider (30):

(30) This surgeon is a butcher.

Clearly, what this utterance evokes is the idea that the surgeon in question is grossly incompetent, dangerous, and so on. The problem, at least for theories of metaphor based on associations or ‘connotations’, is that being incompetent, dangerous, and so on are not properties particularly associated with either butchers or surgeons, so how do these properties emerge when the two categories are associated as in (30)?

If we treat the relationship between an utterance and its interpretation as inferential, then the issue is whether the properties that seem to ‘emerge’ in the metaphorical interpretation can in fact be inferred. It should be obvious that the answer is ‘yes’. Surgeons and butchers both characteristically cut flesh, but in

quite different ways. Surgeons cut live flesh; they cut as little as possible, and with the utmost care to avoid unnecessarily severing blood vessels, nerves or tendons, thus causing irreparable damage. Butchers cut dead flesh to produce pieces of meat for cooking; this places no principled restriction on how much should be cut (or minced, broken, pounded, etc.), and puts a premium on severing nerves, tendons, and other hard tissues. So a surgeon who treats flesh as a butcher does would indeed be grossly incompetent and dangerous. The inferential path to an adequate understanding of (30) involves an evocation of the way butchers treat flesh, and the construction on that basis of an *ad hoc* concept BUTCHER*, denoting people who treat flesh in the way butchers do. Practically all butchers and (one hopes) very few surgeons fall within the extension of this concept. For a butcher, being a BUTCHER* is a quasi-pleonastic property. For a surgeon, on the other hand, it does imply gross incompetence – such an inconceivable degree of incompetence, in fact, that (30) must be seen not just as a metaphor but also as a hyperbole.

A meat lover who cares about precise, careful cuts might praise a butcher by saying:

(31) This butcher is a surgeon.

The interpretation of (31) is symmetrical with the one sketched above for (30), and involves the construction of an *ad hoc* concept SURGEON*, denoting people who cut flesh with extreme care. A butcher who is also a SURGEON* is outstandingly competent and trustworthy. The predicates BUTCHER* and SURGEON*, along with the implication of incompetence for a surgeon who is a BUTCHER* and of competence for a butcher who is a SURGEON*, emerge unproblematically in the course of an inferential comprehension process guided by the search for relevance.

Of course, examples (30)–(31) involve emergent properties that are particularly easy to analyse in inferential terms, and it remains to be seen how far the full range of cases can be dealt with along these lines. However, this account seems promising, and helps to bring out the contrast between inferential approaches to metaphor and more traditional associationist approaches. All inferential relationships are associations, but not all associations are inferential. In claiming that interpretation depends only on inferential relationships, we might have seemed to be depriving ourselves of some explanatory power. As this example suggests, just the opposite is true.¹¹

5.7 Strength of contextual implications, strength of implicatures

We maintain that metaphors are not a distinct category of language use, let alone a discrete one. Are we then denying the obvious truth that metaphors often stand

out as particularly creative and powerful uses of language? If not – and indeed we are not – how are these uses of language to be explained?

Utterances achieve relevance by producing cognitive effects. An utterance may have many cognitive effects or only a few, and these effects may be stronger or weaker. To illustrate, suppose you get to the airport in time for a flight scheduled to arrive in Atlanta at 2 p.m. Hearing an announcement that the flight may be delayed, you say to an airline employee:

(32) I have to be in Atlanta no later than 5 p.m. Will I make it?

She replies as in either (33) or (34):

(33) Your flight will be delayed by at least twenty minutes.

(34) Your flight will be delayed by at least two hours.

Both (33) and (34) imply (35), but only (34) also implies (36):

(35) You have at least twenty minutes to do as you please before boarding.

(36) You have at least two hours to do as you please before boarding.

(35) in turn implies (37), while (36) implies both (37) and (38):

(37) You have time for a drink before boarding.

(38) You have time for a meal before boarding.

Clearly, (34) has more contextual implications than (33).

Both (33) and (34) also provide some evidence for the conclusion in (39):

(39) You will get to Atlanta later than 5 p.m.

Another way of putting this is to say that (33) and (34) *weakly imply* (39). Such weak implications (or probabilifications) are also cognitive effects, and contribute to the relevance of a cognitive input.¹² Since the probability of your arriving late is increased more by (34) than by (33), (39) is a stronger implication (and hence a stronger cognitive effect) of (34) than it is of (33). But still, if you were to assume on the basis of either utterance that you will indeed get to Atlanta later than 5 p.m., this assumption would depend to a considerable extent on your own background beliefs (even more so in the case of (33) than (34)), although it would of course have been encouraged by what the airline employee told you. Overall, this example shows how the contextual implications of an utterance can vary in both quantity and strength.

A competent speaker must have good reason to suppose that what she says will be relevant to the hearer. The hearer himself may provide such a reason: for instance, by asking her a question, thereby letting her know that an answer would be relevant to him. Thus, if a stranger comes up to you in the street and

asks what time it is, you can feel confident that it would be relevant to tell him the time, even if you neither know nor care exactly how it would be relevant, and are implicating nothing more than the presumption of relevance that any utterance conveys about itself.¹³

In most conversations or discourses, the speaker cannot have good reason to think that her utterances will be relevant enough without having some positive idea of the cognitive effects they will achieve. From the hearer's perspective, it is quite often safe to assume that the speaker both expected and intended him to derive some of the implications that he actually does derive, for otherwise she could not reasonably have thought her utterance would be optimally relevant to him. These intended implications are *implicatures* of the utterance. An implicature can be more or less strongly implicated. The speaker may have in mind a specific implication on which the relevance of her utterance depends, and a strong intention that the hearer should derive it; in that case it is strongly implicated. At the other extreme, she may have in mind a vague range of possible implications with roughly similar import, any subset of which would contribute to the relevance of her utterance, and a weak intention, for any of the implications in that range, that the hearer should derive it; these are weak implicatures. Her intentions about the implicatures of her utterance may fall anywhere between these two extremes. The strength of an implicature is determined by the manifest strength of the speaker's intention that a specific implication should be derived. It is important to distinguish the strength of an implicature from the strength of a contextual implication (whether or not it is also implicated), which is the probability that it is true, given that the premise from which it is contextually derived is true.

When the airline employee replies to your question in (32) (about whether you will get to Atlanta by 5 o'clock) as in (33) or (34), she must feel confident that, in telling you how long the delay is likely to be, she is giving you grounds for deriving a weak implication about the risk of your arriving late, thus indirectly answering your question. In other words, you can take her to be implicating that you might indeed be late, leaving it up to you to decide on the seriousness of the risk. The implication is weak – even weaker with (33) than with (34) – but it is fairly strongly implicated.

Does the airline employee also implicate (35) or (36) (viz. that you have at least twenty minutes / two hours to do as you please before boarding)? Although these implications go beyond simply providing an answer to your question, they might help to make the utterance optimally relevant to you in a way the speaker might have both foreseen and intended. When a plane is delayed, people generally want to know how much time they will have to wait before boarding. To that extent, the airline employee may be seen as implicating (35) or (36). These are strong implications of her utterance – they are very probably true – but they are only weakly implicated, because they only marginally increase the relevance of the utterance, and so the speaker's intention to convey them is not

strongly manifest. After all, she may have felt that her utterance was relevant enough without even considering these further implications.

What about (37) (that you have time for a drink), or (38) (that you have time for a meal)? Does the airline employee also implicate these by replying as in (33) or (34)? Again, they are strong implications, which might contribute to the relevance of her utterance in a way the speaker could perhaps have foreseen, but they are even weaker implicatures, because they are among a range of implications with similar import (that you have enough time to buy a magazine, or buy and read one, that you have enough time to do your e-mail, and so on), some of which are likely to be relevant to you although the speaker is not in a position to know which. So she may be encouraging you to consider any of these implications that might be relevant to you, but not any specific one. These are very weak implicatures, if they are implicated at all. By contrast, if your question had been ‘Do I have time for a drink?’ the reply in (33) would strongly implicate (37), and the reply in (34) would strongly implicate (37) and weakly implicate (38), whereas both replies would only weakly implicate (39) (that you will get to Atlanta later than 5 p.m.), if they implicated it at all.

5.8 Poetic effects

Optimal relevance may be achieved by an utterance with a few strong implications, many weak implications, or any combination of weak and strong implications. A speaker aiming at relevance may implicate (that is, anticipate and intend) a few strong implicatures or a wide range of weak implicatures (which may themselves be strong or weak implications). There are many ways of achieving relevance, which differ in both the strength of the implications conveyed and the strength with which they are implicated.¹⁴ Here we are particularly concerned with the case where relevance is achieved through a wide array of weak implications which are themselves weakly implicated. The speaker – or writer, since this method of achieving relevance is particularly well developed in literature – has good reason to suppose that enough of a wide array of potential implications with similar import are true or probably true, although she does not know which these are (hence they are weak implications) and is neither able to anticipate nor particularly concerned about which of them will be considered and accepted by the audience (hence they are weakly implicated). We have argued that the cognitive effects achieved by conveying such a wide range of weak implicatures are identifiable as *poetic effects* (Sperber and Wilson 1995: chapter 4, section 6; Pilkington 2000).

The production of genuinely relevant poetic effects can be a powerfully creative form of language use (creative on the part of both communicator and audience). Such effects can be created by literal, loose, or metaphorical forms of expression. Thus, classical Japanese haikus, which are among the most effective

forms of poetry in world literature, typically involve a literal use of language. Consider Bashō's famous haiku (written in 1680):

On a leafless bough
A crow is perched –
The autumn dusk.

(Translated by Joan Giroux, 1974)

This simple, literal description weakly implicates a wide array of implications which combine to depict a landscape, a season, a moment of the day, a mood, and so on, thereby achieving a powerful overall effect which varies to some extent from reader to reader.

By contrast, many metaphors are not particularly poetic. We are thinking here not so much of conventional metaphors which may have lost their poetic appeal, if they ever had one (was the phrase 'legs of a table' ever poetic?) as of less conventional but not particularly creative metaphors used to highlight a simple idea rather than suggest a complex one. Consider (40), a political comment on the Bush administration's handling of the 2005 Katrina hurricane, compared to its handling of the 2001 terrorist attack on the US:

- (40) Well, if 9/11 is one bookend of the Bush administration, Katrina may be the other. If 9/11 put the wind at President Bush's back, Katrina's put the wind in his face. If the Bush–Cheney team seemed to be the right guys to deal with Osama, they seem exactly the wrong guys to deal with Katrina. (Thomas Friedman, *New York Times*, 7 September 2005)

Here, the use of the metaphors 'bookend' and 'wind' to suggest opposing forces at two ends of a continuum (a case of force dynamics à la Talmy) is so flat that most readers are likely to bypass the obvious relationship between the wind and an explosion, on the one hand, and still more obviously, between the wind and a hurricane, on the other: the cognitive effects derivable from this relationship are unlikely to have been intended, and are hardly worth the effort. Nonetheless, these metaphors serve to make the author's point, which is definitely not of a poetic nature.

Although metaphors are neither necessary nor sufficient for the creation of genuine poetic effects, they are particularly well suited to this purpose, for several reasons. Consider, first, a trivial case of metaphor such as (41):

- (41) WOMAN TO UNCOUTH SUITOR: Keep your paws off me!

Here, 'your paws' refers unproblematically to the hearer's hands. Use of the word 'paws' also activates related notions, conceptions, and images having to do with animal paws, clumsiness, bestiality, and so on. From a relevance theory perspective, the fact that these ideas have been activated suggests that they may be relevant, and the effort spent in activating them, however marginal, suggests that they *should* be relevant (otherwise, the effort would have been wasted,

contrary to the presumption of optimal relevance). While there is a wide range of possible implicatures which might contribute to the relevance of the utterance (that the addressee is clumsy, gross, lusty like a beast, and so on), none of them is strongly implicated by the speaker. We claim that they are weakly implicated: the hearer is indeed encouraged to consider at least some of them and see them as part of the speaker's meaning. It is these vague effects that make the use of 'paws' marginally more relevant than the use of 'hands'.

According to classical rhetoric, the literal meaning of the word 'paw' is replaced in (41) by the figurative meaning HAND. In more recent approaches based on category extension, the literal meaning of 'paw' is extended to include any EXTREMITY OF A LIMB (whether animal or human). In both analyses – substitution of a figurative meaning disjoint from the literal one, or inclusion of the linguistically specified denotation in a broader 'figurative' denotation – suggestions of clumsiness and bestiality are added to the figurative meaning as 'connotations' of the word 'paw'. Here, 'connotations' are associations in a strictly associationist sense: they are grounded in past co-occurrence and can go in any direction.

In fact, the word 'paw' has many associations other than clumsiness and bestiality which might be activated in a metaphor, from the softness of a cat's paw to the strength of a lion's. Association of the type appealed to in associationist psychology is a process which is too vague, on the one hand, and too powerful, on the other, to account for the subtlety and directionality of weak implicatures. As noted above, we would rather appeal only to associations based on properly inferential relationships, and, more generally, stick to an inferential rather than associationist account of comprehension. In (41), the alleged connotations are associated to the literal meaning of 'paw' (i.e. PAW), and not to its figurative meaning HAND or EXTREMITY OF A LIMB. From an inferential point of view, the idea that the literal meaning of 'paw' is discarded while its connotations remain is even more puzzling than the smile of the Cheshire cat: the cat's smile lingers at an empty location, whereas the connotations of the literal meaning of 'paw' are supposed to adorn the figurative meaning that has replaced it.

The alternative analysis we favour is the one we have been defending throughout this chapter. In processing (41), the hearer develops (in parallel) tentative interpretations of the explicit and implicit components of the speaker's meaning, and stops when they fit together in the sense that the explicit content contextually implies the implicated conclusions, and the explicit content and implicit content jointly satisfy the hearer's expectations of relevance. Given that the relationship between explicit content and implicit content is properly inferential, and given the nature of the mutual adjustment process used to determine these contents, the implications evoked by the decoded senses of the words used in the utterance must be genuine implications: that is, they must hold of at least part of the extension of the decoded senses. The *ad hoc* concepts constructed

to carry these implications will then at least overlap with the concepts encoded by the utterance (otherwise we would be dealing with purely associationist rather than inferential relations). Since the concepts PAW and HAND have disjoint extensions, we claim that ‘paw’ in (41) could not be used to convey the meaning HAND. Nor can it be used to convey EXTREMITY OF A LIMB, since this broadened concept is not specific enough to contextually imply clumsiness, bestiality, and so on.

We assume that the *ad hoc* concepts built on the basis of most metaphorical uses of terms are genuinely *ad hoc*: that is, they are adjusted to the precise circumstances of their use, and are therefore unlikely to be paraphrasable by an ordinary language expression. This is why we resort to the ‘*’ notation, and represent the concept pragmatically conveyed by ‘paw’ in (41) as PAW*. PAW* is the most easily constructed concept whose extension includes the hearer’s hands, and which carries the weak contextual implications generally true of prototypical paws: that they are used clumsily, grossly, and so on. These weak implications are themselves weakly implicated: that is, they are weakly intended by the speaker. The utterance on this interpretation achieves optimal relevance by making a strong explicit request that the hearer remove his PAWS*, and weakly implicating that he is behaving clumsily and grossly. Note that PAW*, so construed, involves both a broadening and a narrowing of PAW, as do most *ad hoc* meanings conveyed by metaphorical uses.

So even a common metaphor such as ‘Keep your paws off me!’ achieves some of its relevance through an array of weak implicatures: a poetic touch, however modest. In more creative metaphors, relevance may depend to a much greater extent (or even entirely) on such weak implicatures, in a way that makes it quite appropriate to talk of ‘poetic effects’. Consider the full version of Carl Sandburg’s poem ‘Fog’, whose first two lines are one of the most widely quoted examples of creative metaphor:

The fog comes
on little cat feet.
It sits looking
over harbor and city
on silent haunches
and then moves on.

‘On little cat feet’ evokes an array of implications having to do with silence, smoothness, stealth. Taken together with the following four lines, the phrase evokes a movement which appears both arbitrary and yet composed, so that it is tempting to see it not as random but rather as guided by mysterious dispositions. Poems are read and re-read. On a second reading, the interpretation of the whole poem provides part of the context in which the first two lines are understood. Not unlike Bashō’s literal haiku quoted above, Sandburg’s extended metaphor

weakly implicates an ever-widening array of implications which combine to depict a place, an atmosphere, a mood, achieving a powerful overall effect that varies from reader to reader and reading to reading. It is not part of the explicit content of the poem that the fog comes silently, or smoothly, or stealthily. Rather, what is part of the explicit content is that the fog comes ON-LITTLE-CAT-FEET*. And what is this concept? It is the concept of a property that is difficult or impossible to define, a property possessed in particular by some typical movements of cats (though not all of them – little cat feet can also move in violent or playful ways) and, according to the poem, by the movement of fog. How is this *ad hoc* concept ON-LITTLE-CAT-FEET* arrived at? By taking the poet to be attributing to the coming of the fog that property which contextually implies the very ideas suggested by the phrase ‘little cat feet’.

The example of Sandburg’s poem should help to clarify how and why metaphors are indeed particularly likely to achieve optimal relevance through the creation of poetic effects: the effort required for *ad hoc* concept construction calls for matching effects, and given the freedom left to the interpreter in the construction process, these effects are unlikely to consist in just a few strongly implicated strong implications. It is not that concept construction systematically demands more effort in the case of metaphors (see Gibbs 1994a, Noveck, Bianco and Castry 2001). Many metaphors are very easy to process, while, as any science student knows, arriving at an adequate literal understanding of a statement may take much more effort than a loose or even a metaphorical construal. Nor is it that literal expression is intrinsically less capable than metaphor of achieving poetic effects, as the comparison between Bashō’s haiku and Sandburg’s haiku-like poem shows. It is just that, on the whole, the closer one gets to the metaphor end of the literal–loose–metaphorical continuum, the greater the freedom of interpretation left to hearers or readers, and the more likely it is that relevance will be achieved through a wide array of weak implicatures, i.e. through poetic effects. So when you compare metaphors to other uses of words, you find a bit more of this and a bit less of that, but nothing deserving of a special theory, let alone a grand one.