

# *Metaphor interpretation and emergence\**

ROSA ELENA VEGA MORENO

---

## **Abstract**

A common problem with theories of metaphor is that they do not provide an explanation of what I refer to as the Emergence Problem. Most theories assume that metaphor interpretation consists essentially in mapping some properties of the metaphor vehicle onto the metaphor topic. Thus, in ‘my flatmate is a pig’, the speaker would be understood as attributing to the flatmate some properties of pigs. The problem is that many of the properties which the speaker intends to communicate in using the metaphor are not properties of the vehicle but rather properties which seem to ‘emerge’ in a more global way during comprehension (‘emergent properties’). In this paper I look at the Emergence Problem and, using Relevance Theory, propose a way of solving it.

## **1 The emergence problem**

Early cognitive accounts of metaphor typically assume that understanding a metaphor consists in matching or contrasting properties of metaphor topic and metaphor vehicle so as to identify a subset of properties which they have in common (e.g. Tversky, 1977; Ortony, 1979). Following this line of thought, understanding a metaphor such as ‘my flatmate is a pig’ would involve considering those properties the hearer has stored as part of his knowledge of the speaker’s flatmate and of pigs and selecting a subset of these properties which the speaker’s flatmate and pigs share, for example the properties of ‘being filthy’, ‘being messy’, ‘not being hygienic’, ‘smelling funny’ etc. These properties are taken to form the grounds for interpretation.

Many modern approaches, however, take metaphor interpretation to consist in the attribution of properties of the vehicle to the topic, rather than in a mere match (Glucksberg & Keysar, 1990; Glucksberg, 2001; Gentner, Bowdle, Wolff & Boronat, 2001; Tourangeau & Sternberg, 1981; Tourangeau & Rips, 1991). The hearer of the metaphor above may indeed know certain things about the speaker’s flatmate: he may know that he is English, that he speaks Swahili, that he likes hiking, that he studies law etc. However, he may not know that he is dirty, messy

---

\*I am very grateful to Robyn Carston and Deirdre Wilson for commenting on this paper and for their continuing help and encouragement.

and so difficult to live with. It is this new information that the metaphor aims to provide and that contributes to making the utterance informative to the hearer. Understanding this metaphor involves selecting a set of properties of the vehicle which can apply to the topic, and attributing them to it. Any other property of the vehicle (e.g. pigs have tails) is rejected or suppressed (Gernsbacher et al, 2001). The problem for these models is that sometimes, the set of properties attributed to the topic are not stored as part of our representation of the vehicle, as in (1)-(2).

- (1) Doctor: I am afraid the surgeon who performed a caesarean on your wife perforated both ovaries. I had no choice but to remove them.  
Husband: I want that surgeon out of the hospital. That surgeon is a butcher!
- (2) Jane: I know I have to speak to my boss but I am afraid of him. He is such as bulldozer!

The speaker in (1) may be expressing the thought(s) that his wife's surgeon is highly incompetent, dangerous, careless, etc. The speaker in (2) may be expressing the thought(s) that her boss is stubborn, difficult to deal with, that he is not respectful to her, that he undermines her needs, her thoughts etc. The problem raised by these examples is that our knowledge of butchers does not include the assumption that butchers are negligent and careless and our knowledge of bulldozers does not include the assumption that they are disrespectful or stubborn. Since the set of intended properties are not stored as part of our representation of the vehicle, they can be neither matched with the properties of the topic nor attributed to it. Both matching and attribution models therefore fail to explain how these properties are derived.

Properties which are not part of the hearer's representation for the metaphor vehicle or the metaphor topic, but which seem to emerge in interpreting a metaphor, are often referred to in the literature as 'emergent properties' or 'emergent features'. Examples (1) and (2) show how emergent features play a crucial role in arriving at the meaning the speaker intended to communicate in uttering a metaphor. Any adequate account of metaphor comprehension should be able to provide an explanation for how they are derived. I shall refer to this as the 'Emergence Problem'.

### **1.1 Work on emergence**

Experimental research has shown that 'emergent features' play a fundamental role in metaphor interpretation. Tourangeau & Rips (1991), for instance, found that in providing interpretations for a list of metaphors, subjects produced more emergent features than common features. Furthermore, they judged emergent features as

more relevant to interpretation than either topic-based, vehicle-based or common features. So, for a metaphor such as ‘the eagle is the lion among birds’, a feature such as ‘is respected’, which is associated neither to the topic nor the vehicle individually, was found to figure prominently in subjects’ reported interpretations as well as being judged as highly relevant to those interpretations. Findings like this are repeated across the literature. Gineste & colleagues, for instance, show that over 60% of the properties produced during the processing of poetic metaphors emerge during interpretation (Gineste, Indurkya & Scart, 2000). So, for a metaphor such as ‘the kiss is a fruit’, subjects produced properties such as ‘intense’ or ‘reward’, which are not normally used to characterise either the topic or the vehicle individually. In a series of experiments, Becker (1997) also found that significantly more emergent features and vehicle-based features appear in subjects’ interpretations of metaphors than topic-based or common features. Finally, rather than asking subjects to report interpretations, Tourangeau & Rips (1991) provided subjects with two possible interpretations for a set of poetic metaphors, one based on features common to topic and vehicle, the other based on features which were not commonly associated with either but were nevertheless relevant to interpretation. They found subjects systematically preferred the interpretations based on emergent features.

Scholars generally agree that the existence of emergent properties fits nicely within the interaction view of metaphor (e.g. Gineste et al, 2000). According to Black (Black, 1962), a metaphor such as ‘man is a wolf’ consists of a primary subject ‘man’ (metaphor topic), and a secondary subject ‘wolf’ (metaphor vehicle), each of which is associated to a system of commonplaces which corresponds roughly to the set of assumptions one has about the entities they denote. These include assumptions which are actually true or folk assumptions which, although false, are held as true (e.g. the assumption that wolves are dangerous and aggressive creatures). Metaphor interpretation is seen as resulting from an interaction of such ‘commonplaces’. Black describes the process figuratively.

Suppose I look at the night sky through a piece of heavily smoked glass on which certain lines have been left clear. Then I shall see only the stars that can be made to lie on the lines previously prepared upon the screen, and the stars I do see will be organised by the screen’s structure (Black, 1962: 41)

In understanding the metaphor ‘man is a wolf’, the metaphor topic ‘man’ acts as a frame to highlight commonplaces associated with the vehicle ‘wolf’ (the smoked glass), and the vehicle ‘wolf’ projects back these selected assumptions (the smoked glass with lines on it) which act as a grid to select a set of commonplaces associated with the topic ‘man’ (the set of stars visible through the glass). Looking

at the topic through this grid results in the enhancement of some commonplaces associated to it (visible stars) (e.g. assumptions about men's basic instincts, aggressiveness, competitiveness, etc.), and the suppression of other assumptions (stars which cannot be seen). This reorganisation of assumptions associated with the topic is said to result in the creation of something new, namely a new way of looking at men who are somehow dehumanised.

Black's interactive view of metaphor as a creative process in which two concepts or domains of knowledge interact to form something new has influenced a great deal of psycholinguistic research (e.g. Tourangeau & Sternberg, 1981; Glucksberg & Keysar, 1990; Fauconnier & Turner, 1998; Gineste et al, 2000; Gentner, Bowdle, Wolff & Boronat, 2001). Much of this research sees the emergence of properties during domain interaction as an expected outcome. However, saying that features emerge from interaction is not explanatory: one needs to spell out how it is that they are derived. Being a philosopher, Black's aim was not to provide a detailed account of the pragmatic or cognitive steps involved in interaction or interpretation. Instead, he aimed merely to support the view that metaphors exploit the ability to see something in a new light by means of seeing that thing in terms of something else, and the view that similarities between both entities are created as a result of this interaction. One should then expect the cognitive models inspired by Black's ideas to provide a detailed account of the pragmatic or cognitive steps involved in the derivation of new mental structures and the emergence of new properties. Unfortunately, although a substantial amount of experimental research has been stimulated by the romantic idea of metaphor as powerful and creative, very little work has been done to explain how emergent properties are derived. In fact, experimental work which deals explicitly with the issue, such as that presented above, has mostly been concerned with presenting evidence for the existence of emergent features rather than explanation of the cognitive processes involved in their derivation<sup>1</sup>.

Inspired by the view of metaphor as an interactive process, Glucksberg and colleagues have developed the Class-Inclusion Theory (Glucksberg & Keysar, 1990; Glucksberg, McGlone & Manfredi, 1997; Glucksberg, 2001; McGlone & Manfredi, 2002). According to this theory, metaphor interpretation involves the alignment and interaction of topic and vehicle so as to select a subset of vehicle

---

<sup>1</sup> A possible exception here is the work done in Blending Theory (Fauconnier & Turner, 1998; Grady, Oakley & Coulson, 2000). Followers of this theory propose that the emergence of new structures that result from blending two or more mental spaces often also results in the emergence of new properties which make sense of the blended space. In (Vega-Moreno, forthcoming) I give arguments against this view and propose that the emergence of features during metaphor processing cannot result just from complementing a new mental space with information stored in long term memory (as proposed by Grady et al, 2000); the features need to be inferentially derived.

properties which can attribute a value to a dimension in the topic. This selection of vehicle properties results in the ad hoc construction of an attributive category which the metaphor vehicle is taken to exemplify. These authors claim that metaphors are class inclusion assertions in that they assert that the topic belongs to this new category. Glucksberg and colleagues often illustrate this model with the example ‘my surgeon is a butcher’. They argue that in understanding this metaphor, the hearer aligns vehicle properties and topic dimensions, thus constructing an attributive category ‘people who are incompetent and who grossly botch their jobs’, which the vehicle typifies and which can assign a negative value to the dimension of ‘skill’ provided by the topic (Glucksberg, 2001: 43-55). The problem with this view is that it assumes the attributive category is formed by selecting a subset of properties from the metaphor vehicle (those which can assign a value to a dimension in the topic). However, it is not clear how the property of ‘botching their jobs’ is selected from the vehicle to form the required ad hoc category, when it is a property which is not part of our representation of butchers. Our knowledge of real butchers may include the assumptions that they cut and sell meat, that they use sharp knives, etc. It does not, however, include the assumptions that butchers are incompetent, negligent, careless or people who botch their jobs. If we thought butchers were generally incompetent, we would not trust them and would never buy food from them. Since these properties are not associated to the metaphor vehicle, and since the Class-Inclusion view takes the ad hoc attributive category to be formed by selecting properties from the vehicle, it is not clear how the category of ‘people who are incompetent and who grossly botch their jobs’ is ever formed.

## **2 The transformation problem**

That metaphor interpretation is much more than the selection and attribution of features can be illustrated by a type of emergence problem which I refer to as the ‘transformation problem’, as in (3)-(4).

- (3) Julie: I am afraid about the divorce. My husband’s lawyer is a shark.  
(4) Mary: Are you sure your husband does not mind looking after the children the whole weekend?  
Jane: Yes, don’t worry about it. He is a teddy bear!

The speaker in (3) may want to communicate that her husband’s lawyer is strong and aggressive, that he will attack her in court and persist until he achieves his goals. A case like this presents no apparent problem for attribution theories which take some properties of the metaphor vehicle to be attributed to the topic. In this case, a subset of our knowledge of sharks (e.g. that they are aggressive, persistent,

strong, etc.) is selected in context and attributed to the topic of the metaphor. The Class-Inclusion view often uses the (related) metaphor ‘my lawyer is a shark’ to claim that the hearer takes this metaphor to convey not the assertion that the speaker’s lawyer is an animal which lives in deep waters, but rather the assertion that her lawyer belongs to the category of ‘people and animals who are aggressive and obstinate’ (e.g. Glucksberg, 2001). This category, they claim, is constructed ad hoc by selecting a subset of properties of the vehicle which can assign values to a set of dimensions in the topic. In this case, this would involve selecting the properties of aggressiveness and persistence which can be used to assign a (negative) value to the dimension of ‘skill’ provided by the metaphor topic. One important reason why this account is problematic is that although lawyers and sharks are both aggressive, obstinate and persistent, they are so in very different ways. The property of ‘aggressiveness’ which is attributed to the topic is not the property associated with the encyclopaedic entry of the metaphor vehicle, but a related property which denotes a different kind and degree of aggressiveness. This property, call it aggressiveness\*, seems to ‘emerge’ in interpreting the utterance from this particular subject-predicate combination. Interpreting the metaphor in (3), then, cannot be reduced to the selection of vehicle properties and attribution of these properties; some transformation needs to take place.

(4) presents a clearer case. The speaker of (4) intends to convey some of a range of (weak) implicatures which can be derived by processing her utterance in that context. These include the assumptions that her husband is nice, easy going, always willing to help, easy to please, good with children, etc. If metaphor interpretation involves the attribution of vehicle properties to metaphor topic, the hearer of (4) may access the assumption that teddy bears are soft and cuddly and attribute these properties to Jane’s husband. Unlike (3), in which one can literally say that lawyers can be aggressive, the way in which Jane’s husband is soft is only metaphorical. Thus, understanding the metaphor in (4) cannot be reduced to selecting the physical property of softness common to teddy bears and attributing it to the metaphor topic; some transformation needs to take place.

Although the need to transform the properties of the vehicle into properties that can be appropriately attributed to the topic has been widely acknowledged, very little has been done to provide a solution to the problem. Black himself admits that his model cannot account for it.

A fairly obvious objection to the foregoing sketch of the “interaction view” is that it has to hold that some of the ‘associated commonplaces’ themselves suffer metaphorical change of meaning in the process of transfer from the subsidiary to the principal subject. And these changes, if they occur, can hardly be explained by the account given (Black, 1962: 42)

[B]ecause features are specific to a domain, they must be transformed, i.e. seen in a new way, if we are to find correspondences across domains (Tourangeau & Sterberg, 1981: 217)

The way in which men prey on women is different from the way wolves prey on animals (Gentner & Bowdle, 2001: 227)

The way in which wolves are predators is different from the way men are predators, which in turn is different from the way sharks are predators and lawyers are predators (Gluckberg, 2001: 36)

In an attempt to account for the transformation problem, two main proposals have been considered and rejected. First, the proposal that one may attribute features of the vehicle to the topic on the basis of similarities rather than identity has been rejected on the ground that it would lead to an indefinite regress (Gentner, 1983). Second, the proposal that one may take the assumptions associated to the vehicle as metaphorical has been rejected on the same ground. As Carston points out, considering the encyclopaedic assumptions the concepts encoded by a metaphor give access to as metaphorical “does not break through the metaphorical web” (Carston, 2002a: 87) and so does not allow us to provide an explicit account of how metaphor interpretation takes place. Although I agree that an approach based on similarity of properties cannot adequately account for metaphor interpretation and the transformation of properties, I do think the ‘metaphor within metaphor’ idea is worth exploring. In this paper I want to argue that not only the concepts encoded by the words in an utterance but also some of the concepts that figure in the encyclopaedic entry of these concepts are often pragmatically adjusted during comprehension. This ‘adjustment within adjustment’, which I suggest can shed light on the transformation problem, is not a process unique to metaphor processing but is common to the processing of literal utterances too. In what follows, I will claim that both the transformation problem and the more general emergence problem can be solved using an adequate inferential approach to metaphor interpretation such as Relevance Theory (Sperber & Wilson, 1986/1995).

### **3 Relevance theory and metaphor**

Relevance Theory supports the view that our minds are flexible enough and creative enough to enable us to construct an indefinite range of thoughts, and an indefinite range of concepts which figure as constituents of those thoughts. Following considerable experimental work in Psychology (e.g. Barsalou & Bower, 1980; Barsalou, 1987; Barsalou et al, 1993), Relevance Theory claims that the great

majority of the concepts we form in our minds are unlexicalised concepts which are constructed ad hoc by selecting bits of information from memory and by adjusting concepts that do have a stable entry (Sperber & Wilson, 1998). Most of the concepts and thoughts we form are in fact one-off. They play an important cognitive role, but often have short cognitive lives. Some of these (one-off) concepts and (one-off) thoughts we construct in our minds as we think, we communicate, many others we don't. Some (one-off) concepts and (one-off) thoughts we construct in an attempt to work out what others intend to communicate to us on a certain occasion.

Ostensive communication nicely combines the dynamicity and plasticity of human cognition in constructing new representations with its powerful inferential abilities to recognise the intentions underlying ostensive stimuli. Verbal communication, as a type of ostensive communication, encourages creativity by bringing two minds into the game of recognising each other's intentions, of working out each other's thoughts. Utterance production illustrates the ability of speakers to choose adequate linguistic stimuli with which to communicate to the hearer a range of (lexicalised and unlexicalised) concepts and (one-off) thoughts which they have formed in their minds on a particular occasion. Utterance comprehension involves the ability of hearers to use the concepts encoded by the words the speaker has chosen as clues to those (one off) concepts and (one-off) thoughts she has attempted to convey on that particular occasion. The role of pragmatic theory is to account for how the hearer achieves this.

The pragmatic framework of Relevance Theory offers a psychologically plausible pragmatic approach to utterance interpretation in which pragmatics plays a role not only at sentence level but also at word level (Carston, 1997; 2002a,b; Sperber & Wilson, 1998; Wilson & Sperber, 2002; Wilson, 2004). The approach to lexical pragmatics defended in this framework abandons the code-like assumption that the concept expressed by the use of a word needs to be the very same concept encoded by that word. Instead, in accordance with the view that the stock of concepts we can construct and communicate is greater than the stock of concepts we can encode, it argues that the encoded concept acts merely as a clue to the concept the speaker intended on a particular occasion. The rationale underlying this claim is that if one takes comprehension to be essentially an inferential mind-reading process, there is no reason why the distinction between decoding and inference (which in Relevance Theory is parallel to the distinction between semantics and pragmatics) that operates at sentence level should not also operate at lexical level. That is, just as pragmatic inference bridges the gap between the logical form encoded by an utterance and the set of thoughts communicated in using that utterance, so it should be able to bridge the gap between the concept encoded by the use of a word and the concept the speaker intends to convey in using that word.

Since speakers convey a different concept virtually every time they use a word, a relevance-theoretic pragmatic adjustment process is taken to operate at word level, fine-tuning virtually every encoded concept in context. This process takes as input the concepts encoded by the speaker's utterance, plus contextual information and expectations of relevance, and derives as output the concepts that figure as constituents of the hearer's interpretation of the speaker's thought(s), the ones she intended to convey in using her utterance. It is important to notice that the lexical pragmatic process that operates at word level is not a separate process with its own rules, but rather a special case of the process of mutual adjustment of explicit content, context and implicatures which Relevance Theory takes to operate in interpreting every utterance. It is regulated by the relevance-theoretic comprehension procedure which is automatically triggered in the hearer's mind by any utterance addressed to him, and it takes place in understanding virtually every word, whether it is intended literally, as in (5), approximately, as in (6), hyperbolically, as in (7) or metaphorically, as in (8) and (9).

*Relevance-theoretic Comprehension Procedure*

- a. Follow a path of least effort in computing cognitive effects: Test interpretive hypotheses (disambiguations, reference resolutions, enrichments, implicatures, etc.) in their order of accessibility.
- b. Stop once your expectations of (optimal) relevance are satisfied.

*Communicative Principle of Relevance*

Every utterance communicates a presumption of its own optimal relevance.

*Optimal Relevance*

An utterance is optimally relevant iff

- a. It is relevant enough to be worth the hearer's processing effort.
- b. It is the most relevant one compatible with the speaker's abilities and preferences.

*Narrowing*

- (5) a. The *fish* attacked the swimmer
- b. The *fish* was nice but the potatoes were cold
- c. Please feed the *fish* in my room twice a day

*Broadening*

- (6) I was born with a *square* mark on my foot (approximation)
- (7) You are a *genius*! (hyperbole)
- (8) Getting married and settling down will kill her. She is a *butterfly* (metaphor)
- (9) My daughter, my *princess* (metaphor)

Decoding an utterance such as those in (5)-(9) results in the activation in the hearer's memory of the conceptual addresses encoded by the constituent words in this utterance. These addresses give access in turn to logical properties and encyclopaedic assumptions that are associated to these conceptual addresses (e.g. FISH, BUTTERFLY etc.) in memory. Following a path of least effort in computing cognitive effects, the hearer starts considering these assumptions in their order of accessibility. When an assumption is accessed (e.g. the assumption that fish can be dangerous in (5a), the assumption that butterflies are beautiful and delicate in (8)), it is added to the context and used as a premise to derive implications the speaker might have intended to convey (e.g. the implication that Mary is beautiful and delicate in (8)), which may themselves also be added to the context as premises for further inference. When a contextual implication is derived, the hearer treats it as a potential implicature of the utterance, which may in turn enrich the explicit content via backward inference. This often results in a narrowing or broadening of the encoded concepts. So in (5a) the concept FISH may be narrowed to FISH\*, denoting a subset of fish, namely dangerous fish, and the concept BUTTERFLY may be broadened to BUTTERFLY\* so to denote not just butterflies but beautiful, delicate creatures more generally. If the resulting combination of context, explicatures and implicatures results in an interpretation that satisfies the hearer's expectations of relevance, he stops processing. If not, he considers the next most accessible assumption (e.g. the assumption that butterflies don't stay put in (8)), adds it to the context and tries for further implications. The process of mutual adjustment of explicit content, context and implicatures continues, with assumptions continuing to be considered in their order of accessibility, implications continuing to be derived, explicatures continuing to be enriched, concepts continuing to be adjusted etc., until the hearer arrives at a combination that satisfies his expectations of relevance, at which point he stops.

It follows from this account that although the concepts encoded by an utterance give access in memory to a wide array of assumptions, only a subset of these assumptions is actually processed in context (e.g. the first few assumptions which were most accessible at the moment of processing). The mutual adjustment of explicit content, context and implicatures in order to satisfy the hearer's expectations of relevance often has an effect on these encoded concepts, which are modified accordingly i.e. in the construction of a new ad hoc concept. The denotation of this new concept may be narrower than that of the encoded concept, as in (5), or broader, as in (6)-(8). In fact, one of the greatest contributions of Relevance Theory has been to show that lexical narrowing and broadening are not distinct processes but merely two instantiations of a single process of lexical pragmatic adjustment that fine-tunes virtually every word in context (Carston, 1997, 2002a,b; Sperber & Wilson 1998; Wilson & Sperber 2002; Wilson, 2004). It should not be strange, then, to find cases in which the encoded concept is both

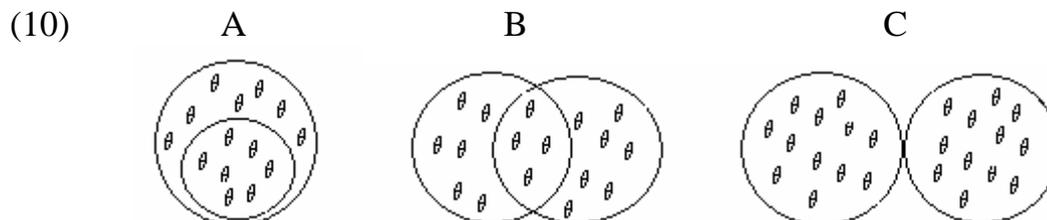
narrowed and broadened during interpretation, as in (9). Arriving at an optimally relevant interpretation after processing a subset of encyclopaedic information associated to the concept PRINCESS (e.g. the assumption that princesses are beautiful, lively and loveable) results in the formation of a new ad hoc concept PRINCESS\*, which is narrower than the encoded concept in that it denotes a subset of real princesses (e.g. princesses who are lively, beautiful and loveable), and broader than the encoded concept in that it denotes a set of entities which typically fall outside its definition (e.g. it denotes a set of young women who, although they are not princesses, are lively, beautiful and loveable).

### **3.1 Pragmatic adjustment and metaphor**

Although all the examples in (6)-(9) involve a process of pragmatic adjustment which results in a broadening of the denotation of the encoded concept, we can distinguish at least two ways in which a concept may be broadened. It may be the case that a word has a strict definition (e.g. exact number, geometric figure etc.) and is typically broadened slightly in use to refer to other cases which (strictly speaking) fall outside that definition; these are often cases of approximation. In (6) for example, the word 'square' is used to refer not only to perfect geometric figures but also to shapes which are roughly square. More radical cases of concept broadening include instances of category extension in which the encoded concept is broadened to include not just a small range of other cases, as in approximation, but a larger range of members which may fall well beyond the scope of the encoded concept. This is the case of words used hyperbolically and metaphorically, as in (7)-(9). In this way, the concept expressed in (7) is broad enough to denote not only real geniuses or near-geniuses but also people who are simply very clever. Similarly, the concept expressed in (8) is broadened to denote people who are beautiful, lively, vulnerable, enjoy freedom and so on.

In Relevance Theory, metaphor interpretation is seen as an instance of pragmatic broadening involving category extension (Carston, 1997, 2002a,b; Sperber & Wilson, 1998; Wilson & Sperber, 2002; Wilson, 2004). In fact, there are at least three possible ways in which the encoded concept can be adjusted in processing a metaphor. The process of pragmatic adjustment of the concept encoded by a metaphorical use of a word may result in the construction of an ad hoc concept which a) denotes all the entities denoted by the encoded concept plus a range of other cases; b) denotes only some of the entities denoted by the encoded concept plus a range of other entities or c) denotes none of the entities denoted by the encoded concept but only a range of other entities. In this last case, the assumptions associated with the encoded concept merely help to construct a new concept whose denotation does not overlap with that of the encoded concept. Providing the circles

stand for sets denoted by concepts, these three possibilities can be illustrated by (10) (see Carston, 2002b: 353).



These three figures can be exemplified by the metaphorical uses in (11)-(12), (13)-(14) and (15)-(16) respectively.

- (11) A. Why does your boyfriend want you to go with him everywhere?  
 B. Because he is a baby  
 BABY\*: denotes people who cannot be independent, cannot look after themselves, can't do things alone, etc. This includes all babies and some adults such as the speaker's boyfriend
- (12) My love, my treasure  
 TREASURE\*: denotes extremely valuable things. This includes all physical treasures and the speaker's love.
- (13) Being the only boy, Dave has always been the prince of the house  
 PRINCE\*: denotes a subset of princes who are spoilt and do as they please, as well as a set of young boys who are not princes but are spoilt and do as they please.
- (14) I am getting divorced because my husband turned out to be an eternal bachelor.  
 BACHELOR\*: denotes a subset of unmarried adult men who party a lot with friends, flirt with women, avoid responsibilities, etc. while excluding others (e.g. the Pope, catholic bishops). It also denotes a set of men who are not bachelors but behave as if they were (e.g. the speaker's husband)
- (15) My boss is a bulldozer  
 BULLDOZER\*: denotes people who are disrespectful, obstinate, undermine people's feelings and thoughts, etc. (e.g. the speaker's boss)
- (16) I tried to persuade Mr Smith to change the essay topic but there was no way. He is an iron bar

IRON BAR\*: denotes people who are difficult to convince, persuade, make change their minds etc. (e.g. the speaker's teacher)

### **3.2 Relevance and emergence**

If we follow the argument above, each of the ad concepts ( $X^*$ ) constructed by the hearer of utterances (11)-(16) would be constructed via a process of lexical pragmatic fine-tuning during the mutual adjustment of explicit content, context and implicatures. This would involve taking, in their order of accessibility, a subset of encyclopaedic assumptions associated to the encoded concepts, which would be added as premises to the context in an attempt to derive an appropriate set of implications. One can see, indeed, how there are stored encyclopaedic assumptions about stereotypical bachelors (e.g. partying a lot, flirting with women, avoiding responsibilities) that might be selected in processing (14) and used to construct an ad hoc concept BACHELOR\* which denotes married and unmarried people who share these properties. It is far more difficult, however, to see how the assumptions that bulldozers\* are stubborn, and that iron bars\* do not have an open mind, which end up being associated to the ad hoc concepts constructed in understanding (15) and (16), are actually selected, since neither of the inanimate objects the encoded concepts denote can have psychological features.

In discussing the related metaphor 'Robert is a bulldozer', Carston (2002b: 350, 2003) suggests that examples like this (in which none of the properties of the metaphor vehicle can apply literally to the metaphor topic) are problematic not only for early and modern psycholinguistic theories of metaphor, but also for the relevance-theoretic approach to metaphor just outlined. Decoding the encoded concept BULLDOZER gives access in memory to a set of assumptions about the physical properties of an inanimate object (e.g. the property of being used in construction, of moving soil etc.). What is not clear is how considering these assumptions, in their order of accessibility, could result in the formation of the concept the speaker is taken to endorse, namely BULLDOZER\*, whose denotation includes a set of obstinate and disrespectful people.

If Relevance Theory claimed that metaphor interpretation consists in constructing an ad hoc concept just by selecting a subset of encyclopaedic assumptions already stored in the encyclopaedic entry of the encoded concept, it would suffer from the same shortcomings as attribution theories (e.g. the Class Inclusion Theory). Like these approaches to metaphor, it would be unable to account for cases of emergence such as those illustrated above. In this paper, I want to argue that this is not the case. Relevance Theory does not assume that metaphor interpretation, or the ad hoc concept construction that takes place during interpretation, reduces to the selection of a set of assumptions associated to the metaphor vehicle in long term memory (and the attribution of these assumptions to

the metaphor topic). Instead, Relevance Theory sees metaphor interpretation as essentially an inferential process, which takes the concept encoded by the metaphorically used word, and the set of assumptions this concept gives access to in memory, as a mere starting point for inference. As will be shown, the gap between the set of existing assumptions associated with the encoded concept and the set of assumptions which end up being associated with the ad hoc concept is bridged by pragmatic inference – thus solving the mystery of emergence. The number of inferential steps taken and the kind of inferential routes followed during the comprehension process results in the formation of concepts whose denotation may depart more or less from the encoded concept, as in (10).

*3.2.1 Emergence and pragmatic inference.* The ‘bulldozer’ and ‘iron bar’ examples in (2), (15) and (16) may be seen as extreme cases of emergence, in that none of the assumptions which end up being associated to the new ad hoc concept are stored in the encyclopaedic entry of the encoded concept. This is partly because the subject of these metaphors is human, while the predicate is linked to inanimate objects. In these examples, virtually every feature, property or assumption associated to the newly formed ad hoc concepts is emergent. We saw that this is not the case for many metaphors in which some of the encyclopaedic assumptions the encoded concept gives access to are shared by the ad hoc concept the speaker intended to convey. In (11), for instance, these may include the assumption that babies require lot of patience, that they cannot do things by themselves, etc. In (14) they may include the assumption that bachelors like partying, that they find it difficult to settle down, that they enjoy freedom, etc. It is worth noticing, though, that even in these cases, the ad hoc concept constructed during metaphor interpretation ends up being associated to a range of other encyclopaedic assumptions about the entities denoted by this new concept, which could not have been retrieved from the encoded concept. In interpreting (11), for instance, the hearer may derive the assumption that the speaker’s boyfriend does not behave in a way suitable for someone of his age. Since babies do indeed behave in a way suitable for people of their age, this is not an assumption the speaker could have retrieved ready-made from his knowledge of babies. In fact, it is an assumption which is stored as part of our representation neither of the speaker’s boyfriend nor of babies, but which seems to emerge in comprehension. Other so-called ‘emergent properties’ which may arise in interpreting this utterance are the property of being spoilt, of being incapable of having a grown-up romantic relationship, etc. (14) is a similar case. In interpreting the metaphor ‘my husband is an eternal bachelor’, the hearer may derive the assumption that the speaker’s husband is not a good husband. This assumption could not have been retrieved ready-made from the encyclopaedic entry of the concept BACHELOR, as bachelors are not married. The property of being a bad husband, which may be attributed to the topic of the metaphor, is another

example of what the literature has referred to as an ‘emergent property’. Other ‘emergent properties’ of the husband which may arise in interpreting this utterance include the property of neglecting his family, of not behaving as expected, of upsetting his wife, of risking his marriage etc. It follows from this that virtually every metaphor enjoys some degree of emergence which needs to be accounted for.

In order to account for emergence, we first need to clarify the nature of an emergent feature. An emergent feature, or an emergent property, has often been defined as a feature/property which is not typically associated to our knowledge of either topic or vehicle, but which arises from their combination. Emergence has been observed not only in the comprehension of metaphor but also in the comprehension of conceptual combinations. Hampton (1997), for instance, observes that in interpreting the combination ‘oxford graduate factory worker’ or ‘rugby player who knits’, people typically produce properties such as ‘failure’ and ‘confused’, respectively, which are not typically associated to any of the terms in the compound. Scholars often distinguish between ‘emergent attributes’ such as those above and ‘extension based emergent attributes’. On some occasions, the properties which arise from the combination of two terms, even when they are not associated to any of the terms individually, seem to be a consequence of identifying a familiar category which the compound can be taken to refer to. So while the property of ‘talking’ is not one we store as part of our knowledge of birds or pets, it is a property which arises in understanding the combination ‘pet bird’ when the compound is taken to denote a subset of pet birds, i.e. parrots. These properties which arise via the identification of an already existing category are commonly referred to as ‘extension based emergent attributes’, in that the identification of the extension of the concept helps to identify the features that characterise the entities denoted (Hampton, 1997; Rips, 1995). It seems to me that a crucial difference between these two types of emergence is that while the property ‘talks’ is retrieved from memory from our knowledge of a type of pet bird, properties such as ‘failure’ and ‘confusion’, which emerge from interpreting combinations such as ‘oxford graduate factory worker’ and ‘rugby player who knits’, can only be inferentially derived. In this paper, I want to claim that many of what the literature has referred to as ‘emergent properties’ are just a set of assumptions which are derived as implications via an inferential process which operates during utterance comprehension. It is because they are implications of an utterance as a whole that they do not need to be associated to any individual term in a combination or a metaphor, but merely need to be derived from the combination of at least two premises used in the comprehension process. Furthermore, because there is a set of inferential steps bridging the gap between the set of assumptions associated to the encoded concepts and the set of assumptions derived as implications of the utterance, it may no longer be appropriate to talk of “emergence”. The following

table illustrates how this inferential process may go for the comprehension of the 'butcher' metaphor in (1)<sup>2</sup>.

(a)	S has said to H 'that surgeon is a butcher'
	<i>Decoding of S's utterance.</i>
(b)	S's utterance is optimally relevant to H
	<i>Expectation raised by the recognition of S's utterance as a communicative act, and acceptance of the presumption of relevance it automatically conveys.</i>
(c)	S's utterance will achieve relevance by elaborating on his immediately preceding comment that he wants the surgeon dismissed
	<i>Expectation raised by (b), together with the fact that such an elaboration would be most relevant to H at this point.</i>
(d)	The fact that a surgeon has operated in a grossly incompetent way is a good reason for wanting him dismissed
	<i>First assumption to occur to H which, together with other appropriate premises such as those below, might satisfy expectation (c).</i>
(e)	A competent surgeon makes incisions in order to preserve life, using high levels of precision, delicacy, foresight and planning to avoid risks
	<i>First accessible assumptions from the encoded concept SURGEON which might combine with (d) and other assumptions to satisfy expectation (c).</i>
(f)	A butcher cuts dead meat in a way that falls far short of the high levels of precision, delicacy, foresight and planning to avoid risk required in a competent surgeon
	<i>First accessible assumptions from the encoded concept BUTCHER which might combine with (d), (e) and a suitably enriched interpretation of (a) to satisfy the expectation in (c).</i>
(g)	The surgeon is a BUTCHER* (where BUTCHER* denotes people who make incisions in a way that falls far short of the levels of precision, delicacy, foresight and planning to avoid risk required in a competent surgeon)
	<i>First enriched interpretation of (a) which might combine with (d), (e) and (f) to satisfy the expectation in (c). Created by pragmatic adjustment of encoded concept by backward inference from the expected conclusion in (h).</i>

<sup>2</sup> I am not claiming here that this is necessarily the sequence in which comprehension occurs. According to Relevance Theory, the mutual adjustment takes place in parallel, rather than in sequence.

(h)	The surgeon operated in a way that falls far short of the high levels of precision, delicacy, foresight and planning to avoid risk required by his job
	<i>Conclusion derived by H from combining (f) and (g). Accepted as an implicature.</i>
(i)	Surgeons who make incisions in a way that falls short of the levels of precision, delicacy, foresight and planning required may cause serious damage to a patient
	<i>Next most accessible assumption from encoded concept SURGEON which might combine with (h) to help satisfy the expectation in (c).</i>
(j)	The surgeon who operated on the speaker's wife caused serious damage to the patient through his lack of precision, delicacy, foresight and planning
	<i>Conclusion derived by H from combining (i) and (h). Accepted as an implicature.</i>
(k)	A surgeon who falls far short of required standards and causes damage to his patient as a result is grossly incompetent
	<i>Next most accessible contextual assumption from encoded concept SURGEON which might combine with (j) to help satisfy the expectation in (c).</i>
(l)	The surgeon who operated on the S's wife was grossly incompetent
	<i>Conclusion derived by H from combining (k) and (j). Accepted as an implicature.</i>
(m)	Grossly incompetent surgeons deserve to be dismissed
	<i>Contextual assumption treated as implicit premise.</i>
(n)	The surgeon who operated on the S's wife deserves to be dismissed
	<i>Conclusion inferred from (m) and (l). Contributing to the satisfaction of (b) and (c). Accepted as an implicature of S's utterance.</i>
(o)	Negligent surgeons are liable to be sued
	<i>Contextual assumption treated as an implicit premise.</i>
(p)	The surgeon who operated on the S's wife is liable to be sued
	<i>Conclusion inferred from (o) and (n). One of several possible weak implicatures of S's utterance.</i>

The husband's utterance, in the circumstances described, raises in the doctor (the hearer) certain expectations of relevance which he expects to satisfy in processing that utterance. At the moment of the utterance, the hearer has certain highly

accessible assumptions, such as the assumption that the speaker's wife has suffered as a result of her operation, that the speaker must be terribly upset about this, that he must be extremely angry with the surgeon and with the hospital, that he and his wife would probably like to make some kind of formal complaint, and so on. In processing the utterance in this context, some of the assumptions associated to the hearer's concept of a butcher become more accessible than others. Following a path of least effort, he starts considering these assumptions in their order of accessibility and adding them to the context in the hope of deriving a set of implications that will satisfy the expectations of relevance raised by the utterance. Because of the presence in memory of the concepts SURGEON and BUTCHER as well as the set of assumptions above, the assumptions that to be a surgeon one needs high levels of precision, delicacy, foresight, etc. and that butchers do not have these qualities may be highly accessible to the hearer at the time. He adds these assumptions to the context and derives the implication that the surgeon fell far short of the required standards for performing his job. This piece of information may trigger further inferences. For instance, combining the information that the surgeon fell far short of the standards required by his job with the assumption that his patient was damaged as a result may lead to the conclusion that he was careless, negligent and liable for sanction, e.g. dismissal or prosecution. The "emergent properties" 'being careless', 'negligent', 'liable to sanction', etc. are thus no more than implications derived inferentially, which would be potentially treated by the hearer as implicatures of the utterance.

It is worth noticing that since utterance interpretation involves a process of mutual adjustment of explicit content, context and implicatures, it follows that as the hearer derives the implications above, the concept conveyed by the word 'butcher' may be continuously adjusted in order to warrant the derivation of the required implicatures. Processing continues with the hearer accessing and combining assumptions, deriving implications and adjusting explicit content until he arrives at a combination of explicit content, context and implicatures that satisfies his expectation of relevance, at which point he stops. As a result of arriving at this combination, the hearer would have constructed a new ad hoc concept BUTCHER\* which denotes the set of people who fall short of the standards of precision, delicacy and foresight required in making an incision, in such a way as to cause damage to living humans and become liable for sanction as a result. It is this concept (or one roughly similar in import) that is taken to be a constituent of the explicature of the speaker's utterance and that warrants the derivation of the implicatures above. Because the encoded concept is merely a starting point for inference, there is no reason why it should not be adjusted to a point where the entities it is normally used to denote fall outside the denotation of the new ad hoc concept that results. The relevance-theoretic account of metaphor, unlike the Class-Inclusion view (and attribution models more generally), can

account not only for emergent properties but also for the comprehension of metaphors falling anywhere on the continuum represented by figures in (10).

I started this paper by presenting the example in (1) as a typical case of emergence, in which understanding a metaphor involves attributing to the topic (i.e. the surgeon) a set of properties which are not generally stored as part of our representation of butchers (e.g. 'falling short of required standards', 'being negligent', 'being careless', 'being liable for damages', etc.). I argued that neither matching models nor attribution models (e.g. the Class Inclusion view) could account for this emergence. We can see now that the reason why modern theories of metaphor cannot provide an explanation for the emergence problem, and so cannot provide a successful account of metaphor interpretation, is partly that they lack an inferential comprehension procedure. Any account of metaphor which sees interpretation as involving simply the selection and attribution of properties without an inferential process operating in between cannot possibly account for metaphor interpretation successfully.

One might argue, however, that even though an inferential approach to metaphor interpretation can account for a subset of emergent properties, it cannot account for how all emergent properties are derived. It might be claimed that it cannot account, for instance, for how a property such as AGGRESSIVENESS\*, denoting a type and degree of aggressiveness typical of people in general and lawyers in particular, is attributed to the topic of the metaphor 'my husband's lawyer is a shark', when the only property the vehicle gives access to is a related, but not identical, property of aggressiveness which does not apply to humans but only to certain animals. Or it might be claimed that an inferential model would not be able to account for how properties such as 'kindness' and 'being easy to please' are derived from our knowledge of teddy bears, or how properties such as 'disrespect' and 'stubbornness' emerge in interpreting the metaphor 'my boss is a bulldozer' in (2). I agree that many inferential approaches to communication would have problems in accounting for these examples, in which some kind of transformation seems to be involved during metaphor processing. However, I want to argue that this is not the case with Relevance Theory.

I have argued that Relevance Theory is built within a picture of cognition which assumes that virtually every concept we form in our minds, both in thinking and in working out what others think, is an unlexicalised concept which is formed ad hoc by selecting different bits of information in memory and adjusting existing concepts. If this is right, then one-off ad hoc concepts figure as constituents of virtually every thought, whether it is explicitly conveyed, implicitly conveyed or not conveyed at all. What I want to point out is that there is then no apparent reason why, in arriving at a hypothesis about the combination of explicit content, context and implicatures the speaker might have intended to convey, pragmatic adjustment should operate only on the set of encoded concepts. Instead, a considerable amount

of pragmatic adjustment may occur narrowing and broadening concepts which are not encoded but which are constituents of thoughts being considered during the interpretation process.

The idea that concepts which are not linguistically encoded but are considered during the inferential process are also adjusted during this process sheds interesting light on the transformation problem. Consider Black's example 'man is a wolf'. In interpreting this utterance, the hearer, following a path of least effort, starts considering assumptions associated to the encoded concept WOLF in the order in which they occur to him. He takes each of these assumptions as a premise and adds it to the context to derive a set of implications that may help to satisfy his expectations of relevance. The expectation that the implications the speaker intended to convey in using this metaphor are consistent with his assumptions about men guides the interpretation and motivates the adjustment of the concepts which figure in the assumptions accessed from the encyclopaedic entry of encoded concept WOLF. A highly accessible assumption the hearer may consider from his knowledge of wolves is that 'wolves are aggressive'. The concept AGGRESSIVE\* as applied to wolves needs to be adjusted on line so that it can warrant the derivation of implications that apply to men.

No metaphor (or utterance) is processed in the absence of a context. Let's consider the possibility that the metaphor above was uttered in a situation where people have been discussing the difficulty of keeping up with a competitive lifestyle. In this situation, the hearer has access to a range of assumptions which he can take as potential implicatures (e.g. the assumption that men are often competitive creatures, that they may undermine others to gain success etc.). These assumptions are used in context to adjust the concept WOLF by backwards inference. In doing so, the concept AGGRESSIVE\* as applied to wolves is adjusted on-line to a point where it warrants the derivation of the expected implicatures (e.g. the implicatures that men are ferociously competitive, may damage others to achieve their own goals, etc.). The same fine-tuning process operates in exploiting some other assumptions about wolves that the hearer may consider, in their order of accessibility, during the comprehension process. He may access the assumption that wolves are predators, which may enable him to derive a range of implications (e.g. they attack other creatures, they only consider their own survival etc.). He takes these implications, together with assumptions about men and business life, as input in order to infer, by mutual adjustment, the ones intended by the speaker. In the process, the concept PREDATOR is adjusted to yield a new concept PREDATOR\*, which applies to men with competitive, aggressive, selfish behaviours, thus warranting the derivation of a set of implicatures which help to satisfy his expectation of relevance (e.g. men are competitive, undermine others to achieve their own success, etc.).

Different expectations of relevance generated by different utterances contribute to a concept being adjusted in different ways. Consider the same property ‘aggressive’, but now attributed to sharks. In processing the metaphor ‘my lawyer is a shark’, and on the assumption that the speaker is happy with his lawyer and confident he is good at his job, the concept AGGRESSIVE may be adjusted to denote a kind of (positive) aggressiveness that involves energy and courage (represented by the ad hoc concept AGGRESSIVE\*\*). However, processing the metaphor on the assumption that the speaker is afraid of his lawyer’s tactics, as in (3), the concept AGGRESSIVE would be adjusted to denote a kind and level of (negative) aggressiveness which involves intentional emotional damage to others (AGGRESSIVE\*\*\*). The concepts AGGRESSIVE\*, AGGRESSIVE\*\* AGGRESSIVE\*\*\* which figure as constituents of our thoughts about sharks in general, and these lawyers in particular, differ from each other, and from the concepts which figure in our thoughts about wolves and men above.

According to Relevance Theory, the ad hoc concepts that result from adjusting the encoded concepts during the interpretation process are taken to be constituents of the explicature of the speaker’s utterance. In this way, the explicatures of the above metaphors would include the concepts WOLF\* and SHARK\* as constituents. These are probably one-off concepts which give access in memory to a set of encyclopaedic assumptions which warrant the implicatures of these particular utterances. What I have tried to show here is that deriving these implicatures has involved a certain amount of pragmatic fine-tuning of other concepts. That is, in constructing the concept intended as a constituent of the explicature of the utterance, other concepts intended as constituents of the implicatures of the utterance are also adjusted. As a result, the encyclopaedic entry of the concept SHARK\*, created on line, would include the assumptions that sharks\* are AGGRESSIVE\* and PERSISTENT\*. The encyclopaedic entry of the ad hoc concept WOLF\* would include the assumptions that wolves\* are AGGRESSIVE\*\* and PREDATORS\*. Since the concepts SHARK\* and WOLF\* are constituents of the explicature of the utterance, adding these assumptions to the context yields implicatures that help to satisfy the hearer’s expectation of relevance (e.g. men are competitive, aggressive, etc.).

It is important to notice that the adjustment of concepts which are not linguistically encoded is not unique to metaphor. Instead, it is a natural by-product of the mutual adjustment process that takes place in understanding every utterance, whether literally, loosely or hyperbolically intended. Consider, for instance, the examples in (17)-(19).

- |      |                      |                                  |
|------|----------------------|----------------------------------|
| (17) | a. The sofa is soft  | Explicature: THE SOFA IS SOFT*   |
|      | b. Baby skin is soft | Explicature: BABY SKIN IS SOFT** |
|      | c. The cat is soft   | Explicature: THE CAT IS SOFT***  |

- |      |                                                                                                          |                                                                                                                                      |
|------|----------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| (18) | a. I love the touch of this sofa<br>b. I love the touch of baby skin<br>c. I love the touch of cat's fur | Implicature: THE SOFA IS SOFT*<br>Implicature: BABY SKIN IS SOFT**<br>Implicature: CAT'S FUR IS SOFT***                              |
| (19) | a. My hair is too long<br><br>b. The cake is ready!<br><br>c. The grass has grown fast                   | Implicature: S WANTS TO CUT* HER HAIR<br><br>Implicature: H CAN NOW CUT** THE CAKE<br><br>Implicature: S WANTS H TO CUT*** THE GRASS |

In (17)-(18), the encoded concept SOFT is adjusted on line to denote a different type of softness on each occasion. In (17), the adjustment of the concept SOFT results in the construction of an ad hoc concept which is taken to be a constituent of the explicature of the utterance. (18) shows how the ad hoc concept formed may be a constituent of one of the premises considered during interpretation as yielding possible implicatures. On some occasions, the (unlexicalised) concept which is a constituent of this premise may have been stored in the encyclopaedic entry of the encoded concept. In other cases, it has to be constructed on-line. (19) presents a similar case in which different ad hoc concepts CUT\*, CUT\*\*, CUT\*\*\* are formed as constituents of a set of assumptions which the speaker might have intended to convey as implicatures of her utterance. In order to arrive at the thought the speaker intended to convey, the hearer needs to adjust a concept which has not been linguistically encoded.

Having looked at the pervasiveness of pragmatic adjustment in utterance interpretation, and at the different inferential steps that may be involved in processing a metaphor, we can now try to see how Relevance Theory can account for problematic metaphorical examples such as 'my boss is a bulldozer', as in (2). In understanding this utterance, the hearer takes both the encoded concept and the assumptions this concept gives access to in memory as input to an inferential process in which he aims to derive the set of assumptions the speaker intended to communicate. Following a path of least effort, he considers potential contextual assumptions about bulldozers in their order of accessibility. These may include the assumption that bulldozers are machines and that they remove obstacles that stand in their way. Although these assumptions cannot be taken to be the ones the speaker intends to convey, they can be used as a starting point to derive hypotheses about those he did intend to convey. That is, they can be used as premises in an inferential process. This process may involve several inferential steps, so that the resulting assumptions may be taken to apply to the speaker's boss.

For example, the hearer may access the assumption that machines are inanimate objects, and that as inanimate objects, they don't have human feelings. Given that

people, including bosses, are expected to have some positive human feelings such as understanding, respect, compassion etc., the hearer may use some of these assumptions to derive potential implicatures. He may think that by portraying him as a bulldozer, the hearer aims to communicate that his boss has no positive human feelings and so is not understanding, respectful or compassionate towards his employees. In order to warrant the derivation of these implicatures, the hearer must adjust the encoded concept BULLDOZER by backward inference into a new concept, BULLDOZER\*, that denotes a set of entities who have no positive feelings for others, who are disrespectful and lack compassion. The hearer may also take the assumptions that the speaker's boss is disrespectful and unsympathetic and combine them with other assumptions which he has stored as part of his knowledge of the relationship between bosses and employees. He can derive from this combination a further array of implications the speaker might have intended to convey as weak implicatures of her utterance (e.g. the speaker's boss sacks people without compassion, he undermines people's ideas, etc.)

Even though the property of being a machine cannot be directly attributed to the topic of the metaphor or be associated to the new ad hoc concept BULLDOZER\*, it can be used as starting point for an inferential process whose output is a set of assumptions which may indeed be accepted as part of the speaker's meaning, and which may end up being stored in the encyclopaedic entry of the new concept. There is at least another inferential route that the hearer may take in processing this metaphor. This will involve the pragmatic adjustment of conceptual material stored in the encyclopaedic entry of the encoded concept BULLDOZER. In processing the utterance 'my boss is a bulldozer', the hearer can access the assumption that bulldozers remove any obstacles that stand in their way. This assumption cannot be directly attributed to the topic. However, it can be broadened to denote a type of removal, a type of obstacles and a range of situations in which obstacles are removed which can warrant the expected set of implications. So, on the assumption that the speaker aims to communicate something about his boss, and that bosses can be difficult to deal with, the hearer may take the encyclopaedic assumption that bulldozers remove obstacles on their way as starting point to construct a broader ad hoc concept [REMOVE OBSTACLES IN THE WAY]\* which denotes the act of despising, rejecting and undermining people, people's feelings and people's thoughts.

A set of other assumptions about bulldozers may be considered during the interpretation process. Many of these assumptions allow the derivation of a range of implications about bulldozers which may be reinterpreted so to warrant a set of implications about the speaker's boss. These may include the assumption that being big heavy pieces of machinery, bulldozers are difficult to move and control. This implication may be reinterpreted in context in the way proposed above to derive the implication that the boss does not have an open mind, that he is inflexible and

stubborn etc. Also, considering the assumption that bulldozers are big, heavy and smash all that comes in their way, one can derive the implication that smaller entities around them are vulnerable and likely to be crushed and destroyed if they stand in their way. These implications may be reinterpreted in context so to apply to the speaker's boss and his relation to his employees. On the assumption that the boss is in a relation of power towards his employees, this allows the hearer to derive a range of implications such as the assumption that the employees are afraid of the boss, of talking to him, of sharing their own thoughts with him; that they feel vulnerable, oppressed and frightened to be reprimanded, humiliated, dismissed etc.

As in the 'butcher' example above, as hypotheses are considered, a continuous adjustment of the encoded concept BULLDOZER takes place so as to warrant the derivation of the expected set of implicatures. This process of mutual adjustment of explicit content, context and implicatures, guided at every point by the speaker's expectations of relevance, and often involving the pragmatic adjustment of concepts that figure as constituents of the assumptions in the encyclopaedic entries of the encoded concepts, continues until the hearer arrives at a combination that satisfies his expectations of relevance, at which point he stops. Providing that the hearer's expectations of relevance are satisfied by the set of implications above, the hearer would have constructed an ad hoc concept BULLDOZER\* which denotes a set of people who are not concerned about people's feelings, who undermine their ideas, thoughts and feelings, who are fixated on their own goals at the expense of others, etc. This concept, which figures as a constituent of the explicature of the speaker's utterance, gives access to an encyclopaedic entry which provides premises for the derivation of implicatures required to satisfy expectations of relevance. These may include the assumption that 'BULLDOZERS\* [REMOVE OBSTACLES IN THEIR WAY]\*'. It is by adding this assumption to the context that a set of (weak) implicatures are derived (e.g. that the speaker's boss undermines her thoughts, her needs, that he does not treat her with respect, that the speaker is afraid he will sack her etc.). The presence of the ad hoc concept BULLDOZER\* as a constituent of the explicature of the speaker's utterance warrants the derivation of these implicatures.

#### **4. Conclusion**

Although the emergence of new properties plays a crucial role in metaphor interpretation, how emergent properties are derived remains still very much a mystery. Using Relevance Theory, I have aimed to offer a way of solving the 'emergence problem' and a version of it which I have labelled the 'transformation problem'. Unlike early and more recent cognitive accounts of metaphor, I have argued that the set of assumptions understood as attributed to the metaphor topic

need not be accessed ready-made from the concepts encoded by the metaphor vehicle. Instead, I have argued that the assumptions made accessible by the metaphor vehicle are merely used as premises in an inferential process which may involve a certain number of inferential steps and a certain amount of pragmatic adjustment before the hearer arrives at an interpretation that satisfies his expectations of relevance. The set of assumptions which result from this process depart (sometimes considerably) from those the encoded concepts gave access to. In other words, what the literature refers to as ‘emergent properties’ are constituents of assumptions derived as implications in processing a metaphor. They do not “emerge” magically in comprehension, but are inferentially derived.

## References

- Barsalou, L. & Bower, B. 1980. A priori determinants of a concept's highly accessible information. Paper presented at the *American Psychological Association Convention*. Montreal.
- Barsalou, L. 1987. The instability of graded structure: implications for the nature of concepts. In U. Neisser (ed.) *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*. 101-140. New York: Cambridge University Press.
- Barsalou, L., Yeh, W., Luka, B.J., Olseth, K.L., Mix, K.S. & Wu, L. 1993. Concepts and Meaning. In K. Beals, G. Cooke, D. Kathman, K.E. McCullough, S. Kita, & D. Testen (eds.). *Chicago Linguistics Society 29: Papers from the parasession on conceptual representations*. 23-61. University of Chicago: Chicago Linguistics Society.
- Becker, A.H. 1997. Emergent and common features influence metaphor interpretation. *Metaphor and Symbolic Activity*, 12/4: 243-259.
- Black, M. 1962. Metaphor. In M. Black, *Models and Metaphors*. 25-47. Ithaca, NY: Cornell University Press.
- Carston, R. 1997. Enrichment and loosening: complementary processes in deriving the proposition expressed? *Linguistische Berichte* 8: 103-127.
- Carston, R. 2002a. Metaphor, ad hoc concepts and word meaning – more questions than answers. *UCL Working Papers in Linguistics*, 14: 83-105.
- Carston, R. 2002b. *Thoughts and Utterances*. Oxford: Blackwell.
- Fauconnier, G. & Turner, M. 1998. Conceptual integration networks. *Cognitive Science*, 22/2: 133-187.
- Gentner, D, Bowdle, B, Wolff, P & Boronat, C. 2001. Metaphor is like analogy. In Gentner, D, Holyoak, K.J, & Kokinov, B.N. (eds.). *The Analogical Mind: Perspectives from Cognitive Science*. 199-253. Cambridge MA: MIT Press.
- Gentner, D. 1983. Structure-mapping: a theoretical framework for analogy. *Cognitive Science* 7:155-170.
- Gentner, D., & Bowdle, B. F. 2001. Convention, form, and figurative language processing. *Metaphor and Symbol*, 16/3-4: 223-247.
- Gernsbacher, M.A., Keysar, B, Robertson, R & Werner, N. 2001. The role of suppression and enhancement in understanding metaphors. *Journal of Memory and Language*, 45: 433-450.
- Gineste, M.D., Indurkya, B. & Scart, V. 2000. Emergence of features in metaphor comprehension. *Metaphor and Symbol* 15/3: 117-135.

- Glucksberg, S. & Keysar, B. 1990. Understanding metaphorical comparisons: Beyond similarity. *Psychological Review*, 97/1: 3-18.
- Glucksberg, S. 2001. *Understanding Figurative Language: From Metaphors to Idioms*. Oxford: Oxford University Press.
- Glucksberg, S. McGlone, M.S & Manfredi, D.A. 1997. Property attribution in metaphor comprehension. *Memory and Language*, 36: 50-67.
- Grady, J., Oakley, T. & Coulson, S. 2000. Blending and metaphor. In G. Steen & R. Gibbs (eds.). *Metaphor in Cognitive Linguistics*. 101-124. Amsterdam: John Benjamins.
- Hampton, J. 1997. Emergent attributes in combined concepts. In T.B. Ward, S. Smith & J. Vaid (eds.). *Creative thought: an investigation of conceptual structures and processes*. 83-110. Washington DC: APA Books.
- McGlone, M.S & Manfredi, D.A. 2002. Topic-vehicle interaction in comprehension. *Memory and Cognition*, 29/8: 1209-1219.
- Ortony, A. 1979. Beyond literal similarity. *Psychological Review*, 86: 161-180
- Rips, L.J. 1995. The current status of research on concept combination. *Mind and Language*, 10, 72-104.
- Sperber, D. & Wilson, D. 1986/1995. *Relevance: Communication and Cognition*. Oxford: Blackwell.
- Sperber, D & Wilson, D. 1998. The mapping between the mental and the public lexicon. In P. Carruthers & J. Boucher (eds.) *Language and Thought*. 184-200. Cambridge: Cambridge University Press.
- Tourangeau, R. & Rips, L. 1991. Interpreting and evaluating metaphors. *Journal of Memory and Language*, 30: 452-472.
- Tourangeau, R. & Sternberg, R. 1981. Aptness in metaphor. *Cognitive Psychology*, 13: 27-55
- Tversky, A. 1977. Features of similarity. *Psychological Review*, 85: 327-352.
- Vega-Moreno, R. Forthcoming. Metaphor: creativity, language and thought. Chapter 2. PhD Thesis. University of London.
- Wilson, D. & Sperber, D. 2002. Truthfulness and relevance. *UCL Working Papers in Linguistics*, 12: 215-254.
- Wilson, Deirdre. 2004. Relevance Theory and lexical pragmatics. *UCL Working Papers in Linguistics* 16