Baltic International Yearbook of Cognition, Logic and Communication

Volume 3 A FIGURE OF SPEECH

Article 10

2007

Metaphor and the 'Emergent Property' Problem: A Relevance-Theoretic Approach

Deirdre Wilson University College London, Department of Linguistics

Robyn Carston University College London, Department of Linguistics

Follow this and additional works at: https://newprairiepress.org/biyclc

080

This work is licensed under a Creative Commons Attribution-Noncommercial-No Derivative Works 4.0 License.

Recommended Citation

Wilson, Deirdre and Carston, Robyn (2007) "Metaphor and the 'Emergent Property' Problem: A Relevance-Theoretic Approach," *Baltic International Yearbook of Cognition, Logic and Communication*: Vol. 3. https://doi.org/10.4148/biyclc.v3i0.23

This Proceeding of the Symposium for Cognition, Logic and Communication is brought to you for free and open access by the Conferences at New Prairie Press. It has been accepted for inclusion in Baltic International Yearbook of Cognition, Logic and Communication by an authorized administrator of New Prairie Press. For more information, please contact cads@k-state.edu.

The Baltic International Yearbook of Cognition, Logic and Communication

Volume 3: A Figure of Speech

August 2008 pages 1-40

DEIRDRE WILSON AND ROBYN CARSTON University College London, Department of Linguistics and University of Oslo, Centre for the Study of Mind in Nature

METAPHOR AND THE 'EMERGENT PROPERTY' PROBLEM: A RELEVANCE-THEORETIC TREATMENT

ABSTRACT: The interpretation of metaphorical utterances often results in the attribution of emergent properties; these are properties which are neither standardly associated with the individual constituents of the utterance in isolation nor derivable by standard rules of semantic composition. For example, an utterance of 'Robert is a bulldozer' may be understood as attributing to Robert such properties as single-mindedness, insistence on having things done in his way, and insensitivity to the opinions/feelings of others, although none of these is included in the encyclopaedic information associated with bulldozers (earth-clearing machines). An adequate pragmatic account of metaphor interpretation must provide an explanation of the processes through which emergent properties are derived. In this paper, we attempt to develop an explicit account of the derivation process couched within the framework of relevance theory. The key features of our account are: (a) metaphorical language use is taken to lie on a continuum with other cases of loose use, including hyperbole; (b) metaphor interpretation is a wholly inferential process, which does not require associative mappings from one domain (e.g. machines) to another (e.g. human beings); (c) the derivation of emergent properties involves no special interpretive mechanisms not required for the interpretation of ordinary, literal utterances.

1. INTRODUCTION: PRAGMATIC ACCOUNTS OF METAPHOR

The goal of a pragmatic account of metaphor is to explain how metaphor is understood, and in particular, how addressees construct an interpretation of the communicator's meaning when a word or other linguistic expression is used metaphorically. This is a special case of the more general pragmatic goal of explaining how addressees bridge the gap between the encoded linguistic meaning of an utterance and the speaker's meaning. Since sentence meaning is often fragmentary and incomplete, and speaker's meaning typically goes beyond it, this gap is pervasive in verbal communication, but it is particularly obvious in cases of metaphorical use. Thus, consider an utterance of (1):

(1) Caroline is a princess.

The linguistically encoded meaning of the word 'princess' is (let's say) the concept PRINCESS, which denotes a subset of female royals. In appropriate circumstances, (1) might be metaphorically used to convey that Caroline, who is not a female royal, is a spoiled, indulged girl, used to special treatment, to having her wishes acted on, to being exempt from the daily chores that others have to perform, and so on. A pragmatic account of metaphor is concerned with how the move from encoded linguistic meaning to metaphorical interpretation is made.

Existing pragmatic accounts differ on several important points. One is their view of how metaphorical use affects the truth-conditional content of utterances (in Grice's terms, what is said; in relevance-theoretic terms, what is explicated). On the standard Gricean account, the speaker in metaphor does not 'say' anything, but merely 'makes as if to say' something that is not itself communicated, but is merely a vehicle for conveying the speaker's implicit meaning, or implicatures. In uttering (1), for instance, the speaker might 'make as if to say' that Caroline is a princess in order to implicate that she is a spoiled, indulged girl (etc.).¹ For a recent defence of this position, see Camp (2006). According to an alternative 'semantic', or 'truth-conditional pragmatic', view, metaphor affects not only the implicatures of an utterance but also its truth-conditional content, and more generally the content of any assertion or other direct speech act that it is used to perform. In (1), for instance, the speaker might be seen as asserting that Caroline is a PRINCESS*, where PRINCESS* is a modification of the encoded concept PRINCESS, and the proposition that Caroline is a PRINCESS* is both a part of what is explicitly communicated and a vehicle for implicatures. This view is held in various guises by Black (1962), Recanati (1995, 2004), Carston (1997, 2002a), Glucksberg, Manfredi, and McGlone (1997),

Sperber and Wilson (1998, 2008), Glucksberg (2001), Wilson and Sperber (2002, 2004), and Wilson and Carston (2007), and we will adopt it here. However, since our main concern in this paper is with the 'emergent property' problem, which arises in all approaches, we will defend the truth-conditional pragmatic approach only where it directly affects the issue of how emergent properties are derived.

A second difference among existing pragmatic accounts is in how closely they are intended to mesh with psycholinguistic investigations of the online comprehension process designed to show, for instance, whether a literal interpretation is always considered before a metaphorical one, or at what stage a particular feature associated with the encoded concept may be activated or suppressed. Let's suppose that the feature FEMALE ROYAL is closely associated with, hence activated by, the encoded concept **PRINCESS**, and is suppressed or inhibited in the course of constructing a metaphorical interpretation of (1). Cross-modal priming experiments might shed light on when (and to what extent) this feature is activated, and when it is discarded or suppressed (see e.g. Gernsbacher, Keysar, Robertson, and Werner 2001; Glucksberg, Newsome, and Goldvarg 2001; Noveck, Bianco, and Castry 2001; Rubio Fernandez 2005, 2007). Theoretical pragmatic accounts of metaphor differ in how far they are intended to be responsive to such findings. Standard Gricean accounts are usually seen as rational reconstructions with few implications for online comprehension, while relevance theory, along with other cognitively oriented approaches, aims to provide an account of metaphor which is not only consistent with existing experimental findings, but itself suggests further experimental tests (see e.g. van der Henst and Sperber 2004).

Existing approaches also differ on whether they treat metaphor as a distinct pragmatic category, or merely as part of a continuum that includes hyperbole, approximation and other local pragmatic phenomena that arise at the level of the word or the phrase. Philosophers of language such as Grice and Lewis seem to have envisaged distinct treatments for metaphor, hyperbole and approximation (e.g. Grice 1967/89: 34, 44-45; Lewis 1975, 1979). Relevance theorists, by contrast, have consistently defended a continuity view, on which there is no clear cutoff point between 'literal' utterances, approximations, hyperboles and metaphors, and they are all interpreted in the same way (for early work, see Sperber and Wilson 1985/6, 1986/95; for a detailed defence of the continuity view, see Sperber and Wilson 2008). The 'emergent property' problem is sometimes raised as a challenge to the continuity view, since metaphorical use is seen as creating emergent properties in a way that non-metaphorical utterances do not (e.g. Romero and Soria 2007). We will argue that the derivation of emergent properties requires no special interpretive mechanisms, and is compatible with a continuity account such as the one proposed in relevance theory.

Finally, existing accounts of metaphor differ in how far they treat metaphor interpretation as properly inferential: that is, as taking a set of premises as input and yielding as output a set of conclusions logically derivable from (or at least warranted by) the premises. At one extreme are predominantly non-inferential, associative approaches, in which PRINCESS in (1), for instance, would be seen as activating, but not implying, associated features such as SPOILED, INDULGED, (etc.). Examples include the computational account proposed by Kintsch (2000) and many treatments of metaphor within the cognitive linguistics framework (Lakoff 1987, 1994; Fauconnier and Turner 2002). At the other extreme are fully inferential approaches such as the one proposed in relevance theory, on which the interpretation of (1) would start from the premise in (2a) and combine it with further contextual premises to derive a conclusion such as (2b):

- (2a) Mary has said 'Caroline is a princess' (where 'Caroline is a princess' is a sentence with a certain—typically fragmentary—decoded meaning, or set of meanings).
- (2b) Mary meant that $Caroline_x$ is a PRINCESS* and a spoiled, indulged girl (etc.).

An intermediate position is taken by Recanati (1995, 2004), who distinguishes 'primary', strictly associative, pragmatic processes from 'secondary', properly inferential, pragmatic processes, with the move from decoded meaning to explicature (e.g. from PRINCESS to PRINCESS*) being treated as a primary, hence non-inferential, process and the move from explicatures to implicatures (e.g. from the premise that Mary said that Caroline_x was a PRINCESS* to the conclusion that Mary meant that Caroline_x was a spoiled, indulged girl (etc.)) as secondary and properly inferential. (On inferential versus non-inferential approaches, see

4

Carston 2002b, 2007; Recanati 2002, 2004; Wilson and Carston 2007; Sperber and Wilson 2008.)

Our main aim in this paper is to argue that the 'emergent property' problem does not present a serious challenge either to the continuity view or to fully inferential accounts of metaphor interpretation. After briefly outlining the relevance-theoretic approach to metaphor in section 2, we will introduce the 'emergent property' problem in section 3, and present our case for a fully inferential treatment of emergent properties in section 4.

2. A RELEVANCE-THEORETIC APPROACH TO METAPHOR UNDERSTANDING

Relevance theory treats metaphor interpretation, like utterance interpretation in general, as guided by expectations of relevance. Relevance is defined as a property of utterances and other inputs to cognitive processes (e.g. external stimuli such as sights and sounds, and internal representations such as thoughts, memories or conclusions of inferences). An input is relevant to an individual when it connects with available contextual assumptions to yield positive cognitive effects (e.g. true contextual implications, warranted strengthenings or revisions of existing assumptions). For present purposes, the most important type of cognitive effect is a contextual implication, which is deducible from input and context together, but from neither input nor context alone. For instance, (3a) might contextually imply (3c) when processed in the context of mentally represented information such as (3b):

- (3a) John lives in London.
- (3b) London is expensive to live in, culturally exciting, with a crumbling infra-structure . . .
- (3c) John has high living expenses, easy access to theatres and cinemas, problems with transport, health care, etc ...

Other things being equal, the greater the cognitive effects, and the smaller the mental effort required to derive them (by representing the input, accessing a context and deriving any contextual implications), the greater the relevance of the input to the individual at that time.

Relevance theory makes two general claims about the role of relevance in cognition and communication. According to the Cognitive Principle of Relevance, human cognition tends to be geared to the maximisation of relevance, so that perceptual, memory retrieval and inferential processes are likely to include automatic heuristics for selecting potentially relevant inputs and processing them in the most relevanceenhancing way. According to the Communicative Principle of Relevance, every act of overt communication conveys a presumption of its own optimal relevance. To be optimally relevant, an utterance (or other act of overt communication) must be at least relevant enough to be worth processing, and moreover the most relevant one compatible with the communicator's abilities and preferences. Together, the Communicative Principle of Relevance and the presumption of optimal relevance ground an inferential comprehension heuristic that provides the basis for deriving a warranted conclusion about the speaker's meaning:

Relevance-theoretic comprehension heuristic

- (a) Follow a path of least effort in constructing an interpretation of the utterance (and in particular in resolving ambiguities and referential indeterminacies, enriching or adjusting the encoded meaning, supplying contextual assumptions, deriving implicatures, etc.).
- (b) Stop when your expectation of relevance is satisfied (or abandoned).

A hearer using this heuristic during online comprehension would proceed in the following way. The goal is to find an overall interpretation that satisfies the presumption of optimal relevance. To achieve this goal, he must enrich the decoded sentence meaning at the explicit level (by disambiguating, assigning reference, and adjusting it in other ways to be discussed below), and complement it at the implicit level (by supplying contextual assumptions which combine with the adjusted explicit meaning to yield enough contextual implications or other cognitive effects to make the utterance relevant in the expected way). What route will he follow in disambiguating, assigning reference, enriching or adjusting the linguistic meaning, constructing a context, deriving contextual implications, and so on? According to the relevance-theoretic comprehension heuristic, he should follow a path of least effort, testing the most accessible referents, disambiguations, contextual assumptions and implications, etc., and stop at the first overall interpretation that yields enough implications (or other cognitive effects) to satisfy his expectations of relevance (see e.g. Sperber and Wilson 1986/95; Carston 2002a; Wilson and Sperber 2004). This is his best hypothesis about the speaker's meaning.

On this approach, any utterance addressed to someone automatically creates a presumption of relevance, which will be satisfied by deriving enough true contextual implications (or other positive cognitive effects), at a low enough processing cost, to make it relevant in the expected way. Given the commitment of relevance theory to a continuity view on which there is no clear cut-off point between metaphorical and non-metaphorical uses, what is true of utterance interpretation in general should also be true of metaphor. And indeed, a metaphorical use of (1) ('Caroline is a princess') might in appropriate circumstances satisfy the hearer's expectations of relevance by contextually implying that Caroline (who is not necessarily a princess) is a spoiled, indulged girl, who expects special treatment, is used to having her wishes granted (etc.). For these implications to be properly warranted, the hearer would have, on the one hand, to construct an appropriate context, and, on the other, to develop the encoded sentence meaning into an appropriate explicit content by disambiguating, assigning reference and enriching or adjusting the linguistic meaning in an appropriate way. These pragmatic processes are seen as taking place not in sequence but in parallel, with tentative hypotheses about context, explicit content and cognitive effects being mutually adjusted or elaborated as online comprehension proceeds. A successful overall interpretation is one that yields enough implications, at a low enough cost, to satisfy the hearer's expectations of relevance, and is internally consistent in the sense that these implications are properly warranted by the context, the presumption of relevance and the enriched explicit content (explicature). (For discussion of this mutual adjustment process, see Sperber and Wilson 1998; Carston 2002a; Wilson and Sperber 2004; Wilson and Carston 2007; Sperber and Wilson 2008.)

According to relevance theory, the explicit content which results from mutual adjustment with context and cognitive effects has typically undergone not only disambiguation and reference assignment, but also modification (or 'modulation') of one or more of the encoded concepts. In (1) ('Caroline is a princess'), for instance, the explicit content might contain not the encoded concept PRINCESS but a related concept PRINCESS*, which is more specific than the encoded concept in some respects and more general in others. The modified concept which is the output of the mutual adjustment process is sometimes called an 'ad hoc' concept, because it is fine-tuned to satisfy the particular expectations of relevance raised by the utterance. (On ad hoc concepts and their contribution to explicitly communicated truth-conditional content, see e.g. Carston 1997, 2002a; Sperber and Wilson 1998, 2008; Wilson and Sperber 2002; Wilson and Carston 2007).

One way of arguing for a continuity view of metaphor is to show that ad hoc concept construction is not specific to metaphor interpretation, but also arises in hyperbole, approximation and even in literal utterances, as a by-product of the same relevance-guided mutual adjustment process. In (4a-c) below, for instance, the ad hoc concept expressed by use of the italicised expression is more specific than the encoded one, and therefore has a narrower denotation:

- (4a) All politicians drink.
- (4b) Buying a house is easy if you've got *money*.
- (4c) No more wine, thanks. I have to get up tomorrow.

Thus, the speaker of (4a) might be understood as asserting not that all politicians drink liquid (an obvious truth), but that they drink alcohol, or, more specifically, that they drink significant amounts of alcohol. Similarly, the speaker of (4b) might be understood as asserting not that buying a house is easy if you have any money at all (an obvious falsehood), but that buying a house is easy if you are suitably rich. Finally, the speaker of (4c) might be understood as asserting not merely that she has to get up at some point or other the next day (which is unlikely to be relevant enough in the circumstances), but that she has to get up early enough for an extra glass of wine to be inadvisable. In each case, the concept the speaker is understood as expressing (i.e. DRINK*, MONEY*, GET UP*) is narrower than the encoded one, applying only to a subset of the items covered by the encoded concepts (DRINK, MONEY, GET UP). In each case, the outcome of the ad hoc concept construction process is an interpretation that would intuitively be classified as literal

(see Wilson and Sperber 2002).

While lexical narrowing happens even in literal utterances, lexical broadening is generally seen as involving some departure from literalness. In (5a-c), for instance, the concepts expressed by use of the italicised expressions might be more general than the encoded ones, with (marginally or substantially) broader denotations:

- (5a) You should take your *empty* bottles for recycling.
- (5b) This policy will *bankrupt* the farmers.
- (5c) The reservoirs are *dry*.
- (5d) The Red Sea is boiling.

In (5a), the word 'empty', which has a relatively strict sense, might be intended and understood as an approximation (involving a relatively marginal broadening of the encoded concept to cover a 'penumbra' of cases which strictly speaking fall outside the linguistically-specified denotation). On this approximate interpretation, the hearer is being urged to recycle not only strictly EMPTY bottles but also bottles which are EMPTY* (i.e. close enough to being EMPTY for the differences to be inconsequential). In (5b), 'bankrupt' may be understood either literally, or as an approximation (BANKRUPT*); it may also be understood as a hyperbole (BANKRUPT**), where hyperbole involves a more substantial broadening of the encoded concept, and hence a greater departure from the encoded meaning. On this interpretation, the speaker of (5b) would be understood as asserting merely that as a result of the policy, the farmers will be substantially poorer than might have been expected or desired. Similarly, in (5c), 'dry' may be used literally, approximately ('almost DRY') or as a hyperbole ('substantially closer to DRY than expected or desired'). Example (5d) illustrates all these possibilities, and one more. As in previous cases, 'boiling' may be understood literally ('at or above boiling point'), as an approximation ('close enough to BOILING for the differences to be inconsequential') or a hyperbole ('closer to BOILING than expected or desired'); it may also be understood metaphorically, as suggesting, for instance, that the water (although not necessarily hot enough to be BOILING, BOILING* or even BOILING**) is bubbling, seething, emitting vapour (etc.).² From this perspective, metaphor interpretation involves a more radical type of broadening than approximation and hyperbole, but, in accordance with the continuity view adopted in relevance theory, arises in essentially the same way. (On the varieties of broadening, see Glucksberg 2001; Wilson 2003; Wilson and Carston 2007; Sperber and Wilson 2008.)

As Carston (1997, 2002a) has shown, narrowing and broadening often combine to vield an adjusted concept that is narrower than the encoded concept in some respects, and broader in others. In (4b) above, for instance ('Buying a house is easy if you've got money'), 'money' might be narrowed, on the one hand, to exclude amounts of money that would be manifestly inadequate in the circumstances, but also broadened to cover not only actual money holdings, but also possessions such as land and art works with a suitable money value. Similarly, in (5a) above ('You should take your empty bottles for recycling') 'empty' might not only be broadened to cover cases where a small amount of liquid is left in the bottle, but also narrowed to cover only a designated type of liquid (e.g. the liquid that the bottle was designed to hold). On this interpretation, the hearer of (5a) is being urged to recycle not only bottles that are strictly or approximately empty, but also, for instance, 'empty' wine or olive oil bottles that have been rinsed and left full of soapy water.³ Metaphor interpretation typically involves a combination of broadening and narrowing. Thus, (1) ('Caroline is a princess'), said of the speaker's younger sister, might be metaphorically understood as expressing a concept PRINCESS* which is broader than the encoded concept in some respects (since it applies to some people who are not actual princesses), and narrower in others (since it applies only to peopleincluding princesses—who are spoiled, indulged (etc.)).

To put a little more flesh on this account, let's assume that the encoded concept PRINCESS is an address or node in memory with three main functions:

- (a) it provides access to mentally represented information about princesses (e.g. the logical information that a princess is necessarily a female royal of a certain type, and a reservoir of more or less strongly evidenced encyclopaedic assumptions about princesses, or particular sub-groups of princesses—that they have public duties to perform, may be spoiled, indulged, etc., and so on)
- (b) it is a constituent of thoughts about princesses, and is therefore activated when thinking about princesses, processing utterances about princesses, etc.

(c) it expresses a property whose extension is the set of (actual or possible) princesses.

When (1) ('Caroline is a princess') is literally understood, it logically implies that Caroline is a female royal, and this affects the truth conditions of the utterance. It does not logically imply that Caroline has public duties to perform, is spoiled, indulged (etc.), since not all princesses have these properties. However, by adding to the context the encyclopaedic information that princesses of a certain type are spoiled, indulged (etc.), and assuming that Caroline is the type of princess of whom these assumptions hold, the hearer may derive the contextual implications that Caroline is spoiled, indulged, (etc.), which may make the utterance relevant in the expected way. The effect of this interpretation would be a narrowing of the encoded meaning PRINCESS to an ad hoc concept PRINCESS*, which denotes only the subset of actual princesses of whom it is true that they are spoiled, indulged (etc.), and a consequent restriction in the content of any assertion the speaker was taken to make. Similar accounts would apply to (4a-c) above.

On this approach, what triggers the narrowing process is the search for relevance (i.e. for enough implications, at a low enough processing cost, to make the utterance relevant as expected). How are these implications derived? By adding to the context encyclopaedic assumptions made accessible by the encoded concept PRINCESS (or by other concepts activated by the utterance or the discourse) and enriching the encoded meaning into an explicit content that combines with these assumptions to yield the expected implications. What direction does the narrowing process take? It follows a path of least effort, considering first the most highly activated contextual assumptions and implications (including those made salient by particular expectations of relevance). When does the narrowing process stop? When enough implications have been derived to make the utterance relevant in the expected way. (For further discussion, see Sperber and Wilson 1998; Carston 2002a; Wilson 2003.)

As expected on the continuity view adopted in relevance theory, a metaphorical interpretation of (1) may be constructed along similar lines, by mutually adjusting context, explicit content and contextual implications so as to satisfy expectations of relevance. Suppose, for instance, that the most obvious referent for 'Caroline' (i.e. the one found

by following a path of least effort in looking for implications) is the speaker's younger sister, who is manifestly not royal. Then the logical information that a princess is a certain type of royal could make no contribution to relevance: its deployment would be a waste of effort, and even if it happened to be automatically activated, it should play no role in interpreting what the speaker meant by uttering (1). The result of dropping this feature would be a concept whose denotation would include all females, and would therefore be considerably broader than the denotation of **PRINCESS**. At the same time, by adding to the context encyclopaedic assumptions made accessible by the encoded concept PRINCESS, and assuming that Caroline belongs to the subset of females of whom these assumptions hold, it may be possible to derive enough contextual implications to make the utterance relevant in the expected way. The effect of adopting these assumptions would be a narrowing of the encoded concept. In these circumstances, the outcome of the adjustment process for 'princess' would be an ad hoc concept PRIN-CESS** which is narrower than the encoded concept in some respects (since it applies only to a subset of actual princesses), but broader in others (since it applies to some people who are not princesses). The resulting overall interpretation might be presented as in (6a-c) (with no constraints on the order in which the premises and conclusions are constructed, and tentative hypotheses about each being mutually adjusted in the course of online comprehension):

- (6a) **Explicit content:**CAROLINE_x IS A PRINCESS^{**}
- (6b) **Contextual assumptions:** A PRINCESS** IS SPOILED, INDULGED (etc.)
- (6c) **Contextual implications:** CAROLINE_x IS SPOILED, INDULGED (etc.)

On this account, both narrowing and broadening are by-products of the search for relevance. What makes (1) intuitively classifiable as 'literal' is the fact that the implications on which the relevance of the utterance depends hold only of actual princesses. What makes (1) intuitively classifiable as an 'approximation', 'hyperbole' or 'metaphor' is the fact that the implications on which the relevance of the utterance depends hold of some things that are not actual princesses (with the difference between 'approximation', 'hyperbole' and 'metaphor' depending on the degree and direction of broadening). In each case, the search for relevance proceeds in the same way, and categories such as 'approximation', 'hyperbole', 'narrowing' or 'broadening' play no role in the interpretation process at all.

To illustrate this point in more detail, consider (7) (a variant of example (5d) above ('The Red Sea is boiling')):

(7) The water is boiling.

As noted above, this utterance might be intended and understood literally, as an approximation, as a hyperbole or as a metaphor, with no clear cut-off point between these possibilities. On the relevance-theoretic account outlined above, all these interpretations are arrived at in the same way: by adding to the context encyclopaedic information made accessible by the encoded concept BOILING (and by other concepts activated by the utterance or the discourse) and deriving enough implications to satisfy the hearer's expectations of relevance. What makes the resulting interpretation intuitively 'literal', 'approximate', 'hyperbolic' or 'metaphorical' is simply the particular set of encyclopaedic assumptions actually deployed in making the utterance relevant in the expected way.

Let's suppose that the encyclopaedic assumptions simultaneously activated by both 'water' and 'boiling' (and therefore potentially highly accessible for the interpretation of (7)) include those in (8a-d):

BOILING WATER: Encyclopaedic assumptions⁴

- (8a) SEETHES AND BUBBLES, HIDDEN UNDERCURRENTS, EMITS VAPOUR, etc.
- (8b) TOO HOT TO WASH ONE'S HANDS IN, TOO HOT TO BATHE IN, etc.
- (8c) SUITABLE FOR MAKING TEA, DANGEROUS TO TOUCH, etc.
- (8d) SAFE TO USE IN STERILISING INSTRUMENTS, etc.⁵

Then (7) would be intuitively 'metaphorical' if the implications that make the utterance relevant in the expected way depend on (8a), but not on (8b-d) (so that the speaker is not understood as committed to the claim that the water is hot)⁶; it would be intuitively a 'hyperbole' if these implications depend on (8b), but not on (8c-d); it would be an 'approximation' if these implications depend on (8c), but not on (8d), and it would be 'literal' if the deployment of (8d) is crucial to making the utterance relevant in the expected way (so that the denotation of

the concept the speaker is taken to have expressed includes only items that are actually BOILING). In each case, the comprehension process works in the same way, by selection of an appropriate set of contextual assumptions to act as premises for the derivation of the expected contextual implications.

The relevance-theoretic account of metaphor comprehension has some similarities to the 'class-inclusion' (or dual reference) account developed by Sam Glucksberg and colleagues (e.g. Glucksberg and Keysar 1990; Glucksberg, Manfredi, and McGlone 1997; Glucksberg 2001). In their view, the word 'jail' in 'My job is a jail' is understood as naming a superordinate category of confining, punishing, inescapable things, which includes actual jails AND the speaker's job. There are several differences of detail between the two approaches, and two more important differences.⁷ First, as already indicated, we locate the account of metaphor within a general account of lexical pragmatic processes of concept modulation or adjustment, which includes both narrowing and several varieties of broadening that would not standardly be treated as metaphorical: that is, we are arguing for a continuity view of metaphor. Second, our aim is to develop a relevance-based account of the cognitive processes that mediate the move from encoded concept to ad hoc concept via mutual adjustment of explicit content, context and contextual implications: that is, we are arguing for a fully inferential account of metaphor. Recently, there has been some evidence of convergence between the two approaches: Glucksberg and colleagues have been exploring the effects of discourse context and considerations of relevance on online metaphor comprehension (Glucksberg 2004), and relevance theorists have been exploring the implications of Glucksberg's work for theoretical pragmatic accounts of metaphor (Rubio Fernandez 2005, 2007; Vega Moreno 2007). Such convergences are likely to benefit research in both psycholinguistics and pragmatics.

3. THE 'EMERGENT PROPERTY' PROBLEM

A certain range of examples has been seen by philosophers, psychologists and pragmatists as presenting a challenge to both continuity and inferential accounts of metaphor. Cases that have been widely discussed include those in (9):

- 15
- (9a) Robert is a bulldozer.
- (9b) Sally is a block of ice.
- (9c) That surgeon is a *butcher*.

In (9a), the speaker might be understood as implicating that Robert is forceful, stubborn, persistent, insensitive to other people's feelings and points of view, and so on; in (9b) she might be understood as implicating that Sally is reserved, unable to express her own feelings, ungenerous or unresponsive to the feelings and overtures of others, and so on; and in (9c), she might be understood as implicating that the surgeon in question is extremely incompetent, dangerous, not to be trusted with the lives of patients, and so on. The question is, where do these implicatures come from? How are they derived?

According to the relevance-theoretic account of metaphor interpretation outlined in section 2, (9a-c) should be interpreted, along similar lines to (6) above, by adding to the context encyclopaedic information made accessible by the encoded concepts BULLDOZER, BLOCK OF ICE or BUTCHER (or by other concepts activated by the utterance or the discourse) and deriving the contextual implications that Robert is forceful, stubborn (etc.), that Sally is reserved, inexpressive (etc.), or that the surgeon is incompetent, dangerous (etc.); if selected by the comprehension heuristic for incorporation into a conclusion about the speaker's meaning these would be accepted not only as implications but as implicatures. But, at the very least, the derivation process cannot be as direct as the one shown in (6) above. In the case of (9a) and (9b), the reason is obvious: our encyclopaedic knowledge of bulldozers, those large machines used for clearing earth, rocks, rubble, etc., is unlikely to include the information that they may be stubborn, persistent, insensitive to the feelings and viewpoints of others (etc.). Similarly, our encyclopaedic knowledge of blocks of ice (solidified H₂O) is unlikely to include the information that they may be reserved, unable to express their own feelings, unresponsive and ungenerous to others (etc.). Only human beings can have psychological properties such as these (taking them literally, as we must, if we care about explanation). (9a) and (9b) are cases of what is often called 'category' crossing: necessary falsehoods, where a literal interpretation of the predicate is incompatible with a literal interpretation of the subject. While (9c) is not a necessary falsehood (although it is factually implausible that the same person would be both a surgeon and a butcher), a similar problem arises, because our encyclopaedic knowledge of butchers is unlikely to include the information that they are incompetent, dangerous, not to be trusted with the lives of patients, and so on.

The question for an inferential account of metaphor is: how can there be a genuinely inferential transition from the premise that the speaker uttered the sentence 'Robert is a bulldozer' to the conclusion that the speaker asserted that Robert_x is a BULLDOZER* and implicated that he is forceful, stubborn, persistent, insensitive to the feelings of others (etc.); and so on for the other examples?⁸ More generally, whatever the proposed account of metaphor, whether it is inferential or not, the question is, how do the properties of forcefulness, stubbornness, insensitivity (etc.) 'emerge' in the course of understanding (9a), when the encoded concept BULLDOZER is literally inapplicable to Robert, and the properties the speaker is understood as attributing to Robert are not listed in the encyclopaedic entry of BULLDOZER (and so on for the other examples)?

These questions have been raised by a number of philosophers interested in metaphor. Pugmire (1998: 99), discussing a metaphorical use of 'iron' similar to the 'bulldozer' case in (9a), comments that 'a predicate does not project unmodified from a non-metaphorical into a metaphorical context. Iron cannot, except metaphorically, be stubborn, persistent, or headstrong'. Martinich (1984/91: 511), considering the possibility that (9b) ('Sally is a block of ice') is understood by supplying the 'commonplace' that blocks of ice are cold and concluding that Sally is cold, notes that an interpretation along these lines 'trades on an equivocation on "cold". In our terms, both these objections make the same point: in order to derive the expected contextual implications, some of the encyclopaedic information associated with the encoded concepts BULLDOZER and BLOCK OF ICE has itself to be metaphorically interpreted, so that the comprehension process involves a metaphor inside a metaphor (or a loose use inside a loose use). While interpretations along these lines seem intuitively plausible in at least some cases, equivocation between premises and conclusion should invalidate an inference, so how is it compatible with a properly inferential account?

In the case of (9c) ('That surgeon is a butcher'), the problem is rather different. It is not so much that some highly accessible information in the encyclopaedic entry for BUTCHER has to be metaphorically interpreted (as 'cold' is metaphorically interpreted in the interpretation of (9b)), but that there is no immediately obvious route from the encyclopaedic entry for BUTCHER to the expected implications at all. As Vega Moreno (2004: 298) puts it,

'Our knowledge of butchers does not include the assumption that butchers are incompetent and dangerous. The properties that the hearer takes the speaker to be attributing to the surgeon are not stored as part of his representation of "butcher", so must be derived by some other means than simply searching through his knowledge about butchers.'

But in that case, how does use of the word 'butcher' contribute to relevance? What encyclopaedic features of butchers, when added to the context and used as premises in inference, would contextually imply that the surgeon in question is incompetent at his job, dangerous to those he is supposed to help, not to be trusted with the lives of patients (etc.)?

The role of emergent features in the comprehension of examples such as (9a-c) has been experimentally investigated (see e.g. Tourangeau and Rips 1991; Becker 1997; Gineste, Indurkhya, and Scart 2000). In a typical off-line experiment, participants are asked to list features they regard as strongly associated with noun phrases in isolation (e.g. 'surgeon', 'butcher'). These noun phrases are then combined into metaphors (e.g. 'That surgeon is a butcher', 'That butcher is a surgeon'), and a different set of participants asked to list the features they take the metaphor to convey. The issue is how far the features listed for a metaphorical utterance as a whole overlap with those independently listed for the metaphor vehicle (i.e. the metaphorically-used predicate), or those common to both metaphor vehicle and metaphor topic (i.e. the subject of the metaphorical utterance), with non-overlapping features classified as 'emergent' (a rather broader conception of emergent properties than the standard philosophical one). The results show that participants tend to cite more emergent features than overlapping features for the metaphor as a whole, and to judge that emergent features are more relevant

to its interpretation than either topic-based, vehicle-based or common features. This is true both for poetic metaphors (Gineste *et al.* 2000) and for more prosaic everyday cases such as (10):

(10) Men are wolves.

In connection with the interpretation of 'wolves' in (10), Tourangeau and Rips (1991: 453) raise a version of the metaphorical reinterpretation problem discussed above in relation to the interpretation of 'bull-dozer' in (9a) and 'block of ice' in (9b). They argue that some of the encyclopaedic properties of wolves must undergo a *transformation* in order to apply appropriately to men; so, for instance, the property of wolves as predators is transformed into the property of competitiveness among men:⁹

'[(10)] is not usually intended to mean that men are carnivorous, although that is a feature shared by men and wolves; instead the metaphor suggests that men are competitive in their dealings with other men, a feature that does *not* characterize wolves.'

Before we go on to consider how the relevance-theoretic approach might account for emergent properties, it is worth emphasising that this is an issue for all pragmatic accounts of metaphor. For predominantly associative (non-inferential) accounts, the question is why the interpretation of metaphors such as (9)-(10) results in the activation of features not activated by the topic or vehicle in isolation. For the standard Gricean account, which treats metaphor as a blatant violation of the first Quality maxim ('Do not say what you believe to be false'), designed to convey a related true implicature, a similar question arises. Assuming that the speaker of (9a) ('Robert is a bulldozer') implicates that Robert ignores the feelings and opinions of others, and that the speaker of (9c) ('That surgeon is a butcher') implicates that the surgeon in question is grossly incompetent, dangerous and not to be trusted with patients' lives, how is the hearer to derive these implicatures on the basis of his encyclopaedic knowledge of bulldozers or butchers, together with other items of background knowledge?

However, the emergent property problem has been seen as presenting a particular challenge to truth-conditional pragmatic accounts of metaphor based on the construction of ad hoc concepts (including the

n mannings may be best se

Metaphor and the 'Emergent Property' Problem

relevance-theoretic account and alternative proposals by Black 1962; Glucksberg 2001; Recanati 1995, 2004). For fully inferential versions of this account (including relevance theory), the challenge is to justify the move from encoded concept to communicated concept, and from communicated concept to implicatures, in examples such as (9)-(10) (see e.g. Carston 2002a; Vega Moreno 2004, 2007). For the continuity view (including the relevance-theoretic version of it outlined in section 2), the challenge is to show that the emergent features of examples such as (9)-(10) can be derived without appeal to special interpretive mechanisms not required for ordinary non-metaphorical utterances. In the next section, we will consider how these challenges might be met.

4. A RELEVANCE-THEORETIC APPROACH TO THE 'EMERGENCE' PROBLEM

The 'emergent property' problem has been recognised and tackled by researchers using the framework broadly known as cognitive linguistics, whose approach to metaphor differs in important ways from the kind of inferential continuity account proposed above. Central to the cognitive linguistics approach is the claim that metaphor is grounded in a system of 'mappings' (i.e. correspondences or associations) between elements from distinct cognitive domains (e.g. the domain of physical properties and the domain of psychological traits, or the domain of machines and the domain of humans) (Lakoff 1987, 1994; Gibbs 1994, 1996; Fauconnier and Turner 1998, 2002). Since no-one, to our knowledge, has suggested that hyperbole or approximation also involve domain mappings (which indeed seems highly unlikely), advocates of a 'mapping' approach to metaphor must reject the continuity view and treat metaphor as a distinct category, with its own special interpretive mechanisms. We see examples such as (7)-(8) above as providing evidence against a 'mapping' account and for the continuity view. The relations between 'domain mapping' accounts of metaphor and fully inferential accounts deserve more detailed exploration than we can give them here, and we hope to address them in future work. For now, we simply note that, if our arguments for the continuity view are correct, and if emergent properties can be derived using only the independently motivated inferential mechanisms outlined in section 2 above, then domain mappings may be best seen as resulting from the repeated use of metaphors bringing together information from the same two domains, and contributing to metaphor interpretation on the effort side, by increasing the accessibility of certain types of contextual assumptions and implications, rather than playing the central role assigned to them in most cognitive linguistic accounts.¹⁰ In the rest of this section, we will suggest inferential analyses of the problematic examples in (9)-(10).

Notice, first, that emergent features are not restricted to metaphor. Several psycholinguistic studies have investigated the derivation of emergent features in intuitively literal conceptual combinations such as those in (11a) and (11b) (see e.g. Rips 1995; Hampton 1997; Glucksberg and Estes 2000):

- (11a) 'casual shirt', 'pet bird', 'digital watch'
- (11b) 'smoky apple', 'sliced tulip', 'upside-down daisy'¹¹

Rips (1995) groups the emergent features of these adjective-noun combinations into two broad types. Some of those associated with the familiar combinations in (11a) are available only to people who happen to have encountered the objects in question. For instance, a 'casual shirt' was described by several participants as one that is PULLED OVER THE HEAD, and a 'digital watch' as one that is RECTANGULAR. It would be hard for someone with no prior experience of casual shirts or digital watches to infer these features on the basis of encyclopaedic knowledge associated with the constituent concepts SHIRT, WATCH, CASUAL and DIGITAL. By contrast, the emergent features of novel combinations such as those in (11b) are inferable on the basis of encyclopaedic knowledge (what Rips calls 'mini-theories') associated with the constituent concepts in isolation. Thus, some participants described a 'smoky apple' as one that TASTES BAD, and an 'upside-down daisy' as one that is UPROOTED. Discussing the 'smoky apple' example, Rips suggests that this second type of feature might be inferred along the following lines:

As a start, our mini-theory for smoky things might specify that they're the result of exposure to heat, usually for an extended period. Our mini-theory of apples is consistent with the possibility that they could be exposed to heat in this way. Furthermore, these mini-theories give us some predictions about the probable effects of this treatment, for instance, that an apple might become dried, hot, blackened, or bad-tasting ... In this way, we can put together a rich composite view, a new mini-theory, of what smoky apples are like that incorporates predictable emergent properties.' (Rips 1995: 100).

Deirdre Wilson and Robyn Carston

Rephrasing this idea in relevance-theoretic terms, we could say that the utterances in (12a) or (12b), when processed in a context made accessible by encyclopaedic knowledge associated with the encoded concepts SMOKY and APPLE, *contextually imply* the conclusions in (12c):

- (12a) This is a *smoky apple*.
- (12b) This apple is *smoky*.
- (12c) This_x Apple is dried, hot, blackened, tastes bad, etc.

Contextual implications are drawn on the individual's own responsibility and not necessarily attributed as part of a speaker's meaning. However, if some of the contextual implications in (12c) are required to make the utterance relevant in the expected way, they would be not only contextual implications but also *implicatures* of (12a) or (12b). Rips was not attempting a pragmatic account of emergent features: he simply presented participants with isolated noun-phrases and asked them to list any features that occurred to them. When a novel conceptual combination—whether literal or metaphorical—is processed in a (real or imagined) discourse context, its interpretation is much more powerfully constrained, and hence much more predictable, for pragmatic reasons.

The presence of a discourse context affects the interpretation of an utterance in two main ways. First, it alters the accessibility of information in the encyclopaedic entries of its constituent concepts, which in turn affects the accessibility of different contextual assumptions and implications. Second, it sets up certain goals or expectations in the hearer. Goal-directed inference is a form of backwards inference from an expected (type of) conclusion to a set of premises that might be used to derive it. As Barsalou (1991) has shown, goal-directed inference speeds up the interpretation process and increases the predictability of the results. So someone processing the phrase 'smoky apple' in the context of the question 'What does a smoky apple taste like?' should find

the emergent feature TASTES BAD doubly easy to derive, by forward inference from activated encyclopaedic information about the effects of food preparation methods on the taste of food, and by backwards inference from the expected type of conclusion A SMOKY APPLE TASTES LIKE

This prediction is confirmed by Glucksberg and Estes (2000), who used a verification task to compare the processing of emergent and nonemergent features (in the broad sense used in psychology) assigned by experimental participants to the conceptual combination 'peeled apple' in different discourse contexts. In one context, the emergent feature WHITE was highly relevant (in both the intuitive sense and our theoretical sense), while in the other, the non-emergent feature ROUND (standardly associated with the constituent concept APPLE) was highly relevant (in both senses), while WHITE was not. The results showed clearly that WHITE was verified faster and more accurately in the contexts where its retrieval made a contribution to relevance, while ROUND was verified faster and more accurately in the contexts where *its* retrieval made a contribution to relevance. Glucksberg (2004: 86) comments:

Apparently, when people understand conceptual combinations in which any number of features are potentially available, feature accessibility is selective, favouring those features that are relevant in the particular context.'

Relevance theory provides a framework in which these effects of discourse context on utterance interpretation can be described and explained. Like most pragmatic theories, it treats utterance interpretation in general as goal-directed. The overall goal is to construct the best hypothesis about the speaker's meaning, and different theories make different proposals about how this is done. According to relevance theory, every utterance addressed to someone creates a presumption of relevance, together with more specific expectations about how relevance is to be achieved (and in particular, about the type of contextual implications to be derived). The hearer's immediate goal is to find an overall interpretation that satisfies these expectations, since this is his best hypothesis about the speaker's meaning. As described above, the relevance-theoretic comprehension heuristic is an automatic inferential procedure for constructing such an interpretation by following a path of least effort in mutually adjusting context, explicit content and contextual implications (via both forward and backward inference) so as to make the utterance relevant in the expected way. Implicatures are contextual assumptions and implications that have to be added to the interpretation in order to satisfy the expectations of relevance raised by the utterance. We will try to show that these strong pragmatic constraints on inferential comprehension play a central role in the derivation of emergent features in metaphorical utterances, including the problematic examples in (9)-(10) above.

Returning to our original example in (1) ('Caroline is a princess'), let's consider how it might be understood in the discourse context in (13a), a question about the addressee's younger sister Caroline, who is manifestly not a princess:

- (13a) Will Caroline help us clear up the flood damage?
- (13b) Caroline is a princess.

The hearer's goal in interpreting (13b) is to derive an answer to his question (i.e. a conclusion of the form $CAROLINE_x$ WILL/WON'T CLEAR UP THE FLOOD DAMAGE). This could be done by enriching the encoded sentence meaning as in (14a) and supplying the contextual assumption in (14b):

- (14a) **Explicit content:** CAROLINE_{*x*} IS A PRINCESS^{*}
- (14b) **Contextual assumption:** A princess* doesn't clear up flood damage.
- (14c) Contextual implication: $CAROLINE_x$ WON'T HELP US CLEAR UP THE FLOOD DAMAGE.

Of course, the contextual assumption in (14b) is unlikely to be stored ready-made in the encyclopaedic entry for PRINCESS, and to that extent the interpretation of (13b) involves the derivation of an emergent feature (DOESN'T CLEAR UP FLOOD DAMAGE) in the broad sense used by psychologists. However, this feature would be straightforwardly derivable in the course of the mutual adjustment process, by a combination of forward inference from existing encyclopaedic features (e.g. UNUSED TO PERFORMING MENIAL TASKS, UNACCUSTOMED TO MANUAL LABOUR), and backward inference based on the expected type of conclusion in (14c). What justifies the choice of this interpretation over alternative, logically

possible ones (e.g. interpretations suggesting that Caroline will help clear up the flood damage) is the fact that this is the first accessible interpretation to make the utterance relevant in the expected way, and it is therefore the one selected by the relevance-based comprehension heuristic. Thus (14c) would be accepted not only as a contextual implication but also as an implicature of the utterance in (13b).

Vega Moreno (2004, 2007) has argued that the emergent features of (9c) ('That surgeon is a butcher') can be inferentially derived along similar lines. Here is a slightly adapted version of her account of how the derivation might go. Suppose someone utters 'That surgeon is a butcher' immediately after uttering (15a):

- (15a) That surgeon ought to be dismissed.
- (15b) He is a butcher.

The processing of (15a) would activate the hearer's knowledge of surgeons (and of the particular surgeon referred to in the utterance), which might include the logical feature IS A DOCTOR and more or less strongly evidenced encyclopaedic assumptions such as those in (16): SURGEON:

Encyclopaedic assumptions

- (16a) WORKS IN A HOSPITAL, IN STERILE CONDITIONS, etc.
- (16b) OPERATES ON HUMANS WITH CONCERN FOR THEIR WELFARE, etc.
- (16c) CUTS FLESH LIKE THIS: [XXX], etc.
- (16d) REQUIRES GREAT DEXTERITY, MEDICAL TRAINING, EDUCATION, etc.

Here, '[xxx]' is meant to stand for a representation (conceptual, sensory or kinaesthetic) of a surgeon's manner of cutting flesh, broadly construed to include information about the techniques and instruments used, skills deployed, underlying intentions, physical conditions and consequences, and so on.¹² The processing of (15a) is also likely to raise a question in the hearer's mind about why the speaker thinks the surgeon ought to be dismissed (is it for negligence or incompetence, for moral turpitude, for quarrelling with his colleagues, as a cost-cutting measure, and so on?), and an expectation that the next part of the utterance will answer it by conveying a conclusion of the form THAT SURGEON_x OUGHT TO BE DISMISSED BECAUSE ____.

The processing of (15b) would activate the hearer's knowledge of butchers, which might include the logical features IS A TRADESMAN, SELLS MEAT and more or less strongly evidenced encyclopaedic assumptions such as those in (17):

BUTCHER: Encyclopaedic assumptions

- (17a) WORKS IN A SHOP; SELLS BEEF, LAMB, PORK, POULTRY, etc.
- (17b) CUTS UP DEAD BODIES FOR USE IN COOKING, etc.
- (17c) CUTS FLESH LIKE THIS: [YYY], etc.
- (17d) REQUIRES VOCATIONAL TRAINING AND SKILLS, etc.

Here, assumptions (17b) and (17c) would be primed by the prior mention of 'surgeon' (given that a surgeon also cuts flesh). But a surgeon who cuts flesh in the way a butcher does (using the same techniques, with the same intentions, concern for welfare, degree of skill, etc.) would be grossly incompetent and dangerous to patients, and would deserve to be dismissed. In the course of the mutual adjustment process, it would therefore be relatively easy to construct an overall interpretation on which the speaker of (15) is understood as asserting that the surgeon in question is a BUTCHER* (where a BUTCHER* is a person who cuts flesh in a way appropriate to butchers), and implicating that he ought to be dismissed because, being a BUTCHER*, he performs operations in a grossly incompetent, dangerous way. This account of the derivation of the emergent features of (9c) is genuinely inferential, with no appeal to special interpretive mechanisms such as domain mappings.¹³

This account is oversimplified in one obvious respect. Although (15) may evoke images of a surgeon hacking at flesh in the way a butcher does, and these may put the hearer on the track of an overall interpretation that would satisfy his expectations of relevance by explaining why the surgeon was dismissed, the assumption that the speaker meant that the surgeon was a BUTCHER* in this sense (i.e. that he cuts flesh in just the way a butcher does) is both factually and pragmatically implausible, and is unlikely to be accepted as it stands. To put the same point another way, 'butcher' in (15) is not only a metaphor but a hyperbolic metaphor: BUTCHER* suggests a satisfactory line of interpretation, but has to be broadened still further, retaining only those features of the way a butcher cuts flesh that can also be plausibly attributed to a

few surgeons and would help to explain why the surgeon in question was dismissed. Thus, the speaker of (15) might be understood as asserting that the surgeon in question was a BUTCHER** (i.e. that he cuts flesh with a degree of skill, delicacy, preparation, reflection and concern for the physical consequences of his actions that are plausibly attributed to butchers and a few surgeons), and implicating that, being a BUTCHER** he operates in a grossly incompetent way and deserves to be dismissed.¹⁴ Here, 'operating with a certain degree of skill, etc.', 'operating in a grossly incompetent way, etc.' and 'deserving to be dismissed, etc.' are all emergent features in the following sense: they are not encyclopaedic features of BUTCHER, BUTCHER* or even BUTCHER** (since they do not apply to actual butchers), but are inferentially derivable as contextual implications of the assertion that the surgeon is a BUTCHER**, given a context containing standard encyclopaedic information about surgeons.

Similar accounts can be given for at least some of the category crossing cases in section 3. Suppose that (10) ('Men are wolves') is uttered during a conversation about how two business partners, Smithers and McGee, have been trying to defraud each other and take over the business profits. Like (15), (10) is both a metaphor and a hyperbole. The simplest account of the interpretation process might go as follows. The utterance of (10) would activate the encoded concept wolf, with (let's say) the logical feature ANIMAL OF A CERTAIN KIND and more or less strongly evidenced encyclopaedic assumptions such as those in (18) (many no doubt based on cultural stereotypes rather than biological knowledge):

WOLF: Encyclopaedic assumptions

- (18a) BY NATURE AGGRESSIVE, VICIOUS, MERCILESS, PREDATORY, SAVAGE
- (18b) SOLITARY OR HUNTS IN PACKS; HOSTILE TO HUMANS, FRIGHTENING

Several of these features also apply to (some) men, and would be simultaneously activated by the discourse context and expectations of relevance. As a result of the mutual adjustment process, the speaker of (10) might therefore be understood as asserting that men in general are WOLVES* (where a WOLF* is by nature aggressive, vicious, merciless, savage, solitary, hostile to humans, etc.), and implicating that aggressive, vicious, merciless, savage (etc.) behaviour is only to be expected

In fact, as Tourangeau and Rips (1991) point out, the features that actually figure in our encyclopaedic entry for WOLF may be more specific than those suggested above, because wolves are vicious, aggressive, savage and merciless (etc.) in a particular, wolf-like way. If so, these more specific features might be represented as narrowed concepts (e.g. VICIOUS^{*}, SAVAGE^{*}, along the lines discussed above for examples (4a-c)), where vicious* is paraphraseable as vicious in the wolf-like way, and so on for the other features. Each of these narrowed concepts would inherit many of its encyclopaedic features from the more general concept (VICIOUS, SAVAGE, etc.) from which it is derived, but it would also have a range of further features that apply specifically to wolves. In interpreting (10), the hearer's goal is to find a set of encyclopaedic features of VICIOUS*, SAVAGE*, etc. which apply not only to wolves but also to some humans, and would explain why Smithers and McGee have been trying to defraud each other even though they are business partners. Such features might include, for instance, being vicious, savage, etc. even to one's own kind, and in a particularly intense, instinctive, unreflective, physically aggressive, reckless and inhumane way. In that case, the speaker of (10) might be understood as asserting (as before) that men are wolves* and implicating that Smithers and McGee, being wolves*, are by nature vicious, savage etc., even to each other, and in a particularly intense, instinctive, unreflective, physically aggressive, reckless, inhumane way.¹⁵ Either account is compatible with a fully inferential treatment of metaphor interpretation.

Let's now look briefly at (9b) ('Sally is a block of ice'), recalling Martinich's comment (see section 3 above) that inferential accounts of this example run the risk of being invalidated by an equivocation in the understanding of 'cold'. To meet this challenge, we have to show how, from the premise that the speaker has uttered the sentence 'Sally is a block of ice', together with other easily accessible contextual assumptions, the hearer can validly infer that Sally is emotionally reserved, unaffectionate, unresponsive to the overtures of others, etc.. At least part of the account is straightforward, and the interpretation proceeds along similar lines to the ones sketched above for (9c) ('That surgeon is a butcher') and (10) ('Men are wolves'). Suppose that Jenny utters 'She is a block of ice' immediately after uttering (19a):

- (19a) I had dinner with Sally last night.
- (19b) She's a block of ice.

This utterance would automatically activate the encoded concept BLOCK OF ICE, which has (let's say) the associated logical feature FROZEN WATER OF A CERTAIN FORM and more or less strongly evidenced encyclopaedic assumptions such as those in (20):

BLOCK OF ICE: Encyclopaedic assumptions

- (20a) SQUARE, SOLID, HARD, COLD, RIGID, INFLEXIBLE, etc.
- (20b) DIFFICULT/UNPLEASANT TO TOUCH, COME CLOSE TO, INTERACT WITH, etc.
- (20c) MAKES THE SURROUNDING ATMOSPHERE UNCOMFORTABLE, etc.
- (20d) MAKES PEOPLE WANT TO MOVE AWAY, etc.

Several of these encyclopaedic features apply straightforwardly to (some) humans as well as blocks of ice, and might also be activated by the discourse context and expectations of relevance. At the same time, the logical feature FROZEN WATER provides access to a 'mini-theory' of how water is transformed by the freezing process from a natural substance which is soft, flexible and adapts to its surroundings, to a hard, rigid, inanimate-seeming object which is incapable of adapting to its surroundings. As a result of the mutual adjustment process, Jenny might thus be understood as asserting that Sally is a BLOCK OF ICE* (where a BLOCK OF ICE* is hard, rigid, inanimate-seeming, difficult to interact with, unpleasant to touch or come close to and incapable of adapting to its surroundings), and implicating that her evening with Sally was not a success because, being a BLOCK OF ICE*, Sally has a limited capacity for human interaction or responses such as conversation and the expression of emotion.

So far, we have only considered features of the encoded concept BLOCK OF ICE which apply straightforwardly to humans. However, we assume, as pointed out by Martinich and others, that physical descriptions such as 'hard', 'rigid', 'inflexible', 'square', 'solid', 'cold', 'icy', 'frozen', etc., apply to humans only in an extended non-physical sense, which is now presumably lexicalised but which arose via metaphorical extension of the basic physical sense. The question is how the hearer can get from the basic physical concepts which (presumably) feature in (20) above to broader senses that can be appropriately applied to Sally. According to Martinich, there is no genuinely inferential route.

We want to argue that, given the relevance-theoretic account of metaphor outlined above, there are in fact two possible inferential routes, each of which is likely to be exploited in at least some cases. The first route involves taking seriously Martinich's suggestion that the psychological senses of 'hard', 'rigid', 'cold', etc. are metaphorical extensions of the basic physical senses. If so, it follows from the relevance-theoretic analysis of metaphor that these extended senses arose through repeated broadening of the basic physical senses (HARD, RIGID, COLD, etc.) to create superordinate concepts (HARD*, RIGID*, COLD*, etc.) which are not purely psychological but have both physical and psychological instances. On this approach, the denotation of the basic, physical sense would be partially included in the denotation of the broader superordinate sense, and the encyclopaedic features of the superordinate sense and the basic physical sense would overlap. If so, then a block of ice can be truly described not only as HARD, RIGID, COLD, etc, but also as HARD*, RIGID*, COLD*, etc, and both the basic physical features and the more general physical/psychological features would figure in the encyclopaedic entry for BLOCK OF ICE. For a hearer following this inferential route, the speaker of (9b) ('Sally is a block of ice') or (19b) ('She's a block of ice') might be understood as asserting that Sally is a BLOCK OF ICE* (where BLOCK OF ICE* applies both to blocks of ice and to some humans), and deriving the contextual implications that Sally is HARD*, COLD*, RIGID* (etc.) (where HARD*, COLD*, RIGID* etc. are applicable both to humans and to blocks of ice, and would be highly activated by the discourse context and expectations of relevance). Of course, HARD*, RIGID*, COLD* have their own encyclopaedic entries, some of the contents of which (e.g. RESERVED, UNRESPONSIVE TO THE FEELINGS OF OTHERS, UNGENEROUS, etc.) apply only to the subset of humans in their domain and might interact with contextual information about female humans (and Sally in particular) to yield further contextual implications. As a result of the mutual adjustment process, the speaker of (19b) might therefore be understood as asserting that Sally is a BLOCK OF ICE*, and implicating that she did not enjoy dinner with Sally because, being a BLOCK OF ICE*, Sally is HARD*, COLD*, RIGID* (etc.), and therefore emotionally reserved, unresponsive to other people, unable to express her

own feelings, unpleasant to interact with, be close to, and so on.

In a study of a wide range of these 'double-function' adjectives, Asch (1955; 1958) explores the idea that there is a unitary conceptual basis to the use of 'cold', 'hard', etc. to describe both physical and psychological properties. In his view, these inclusive concepts are grounded in our observations of and interactions with our fellow humans, aspects of whose behaviour and appearance we experience as relevantly similar to our experience of interacting with physically cold/hard/etc. objects:

'The hardness of a table and of a person concerns events radically different in content and complexity, but the schema of interaction is experienced as dynamically similar, having to do with the application of force and of resulting action in line with or contrary to it. What holds in the preceding instance applies to the other terms in the same category. *Warm*, aside from thermal qualities, stands for bringing closer, or for drawing into a union, while *cold* excludes or isolates.' Asch (1958: 93)

Recast in our terms, what Asch is suggesting is that there is a lexicalised superordinate concept (COLD*, HARD*, RIGID*) which applies both to objects that we find cold to the touch and to people whose personality we would describe as cold, and which would be deployed in the interpretation of 'Sally is a block of ice'.¹⁶ This fits with our analysis of how 'Sally is a block of ice' would be understood by a hearer following the first of our two possible inferential routes.

A second possible inferential route would start from the assumption that polysemous words such as 'hard', 'rigid', 'cold', etc. have distinct lexicalised physical and psychological senses (HARD and HARD**, RIGID and RIGID**, COLD and COLD** etc.), whose denotations do not overlap in the way described above for COLD and COLD*, HARD and HARD*, (etc.). On this analysis, these psychological senses would not be metaphorical extensions of the basic physical senses (although they might have arisen, in the history of the language or the individual, via narrowing of such broader superordinate senses). Still, in our framework, where non-lexicalised ad hoc concepts may be constructed on the fly in order to satisfy expectations of relevance, it is easy to see how ad hoc (nonlexicalised) superordinate concepts such as COLD*, HARD*, (etc.), whose denotations include both items that are COLD/HARD and items that are COLD**/HARD**, might be constructed during the online interpretation of (9b) ('Sally is a block of ice') or (19b) ('She's a block of ice'). In this case, the superordinate concepts COLD*, HARD*, RIGID*, etc. would be broadenings of the basic physical concepts COLD, HARD, RIGID, (etc.), with overlapping encyclopaedic features and denotations. As a result, BLOCK OF ICE would contextually imply COLD*, RIGID*, HARD* (etc.), and the interpretation would proceed as for the first inferential route described above.

Independent evidence that at least one of these inferential routes must be not only available but exploited in utterance interpretation comes from similes such as (21), where 'cold' must be understood as expressing a concept that is general enough to apply simultaneously to both psychological and physical objects:

- (21a) I had dinner with Sally last night.
- (21b) She's as cold as a block of ice.

Attested examples of such similes from the British National Corpus include those in (22):

- (22a) His mind was as cold as the ice forming on the windscreen.
- (22b) His voice was as cold as the Arctic snows.
- (22c) His own voice was low and as cold as steel.
- (22d) His eyes were as cold as polar ice.
- (22e) His silvery-green eyes looked as cold as glacial ice.
- (22f) He's good and great, but as cold as ice.

These do not seem to involve a pun or equivocation on 'cold'. If so, they provide further evidence that hearers are capable of accessing and using a superordinate concept COLD* (whether lexicalised or non-lexicalised) whose denotation includes both physical and psychological instances and hence support our inferential account of how the emergent properties of (9b) ('Sally is a block of ice') are derived. (We are indebted to Dan Sperber for discussion of this kind of example.)¹⁷

These inferential routes to the derivation of emergent properties apply equally to (9a) ('Robert is a bulldozer'). The metaphorical use of 'bulldozer' has many possible interpretations, some more concerned with physical appearances (Robert is physically big, heavy, clumsy in his movements, pushes people aside in order to get past, etc.), others more focused on psychological characteristics (Robert is a forceful personality, unstoppable when he has decided on a course of action or is pursuing an idea, etc.).¹⁸ Suppose that two members of a university department, Mary and Robert, have very different views on how to cope with the recent announcement that their department's funding is to be severely cut. Mary is reluctant to discuss her ideas with Robert, commenting, 'Robert is a bulldozer'. In this discourse context, (9a) would activate the encoded concept BULLDOZER, with (let's say) the logical feature MACHINE OF A CERTAIN KIND and more or less strongly evidenced encyclopaedic assumptions such as those in (23):

BULLDOZER: Encyclopaedic assumptions

- (23a) LARGE; POWERFUL; CRUSHING; DANGEROUS TO BYSTANDERS, etc.
- (23b) LOOKS LIKE THIS: [XXX]; MOVES LIKE THIS: [YYY], etc.
- (23c) GOES STRAIGHT AHEAD REGARDLESS OF OBSTACLES, etc.
- (23d) PUSHES ASIDE OBSTRUCTIONS; DESTROYS EVERYTHING IN ITS PATH, etc.
- (23e) HARD TO STOP/RESIST FROM OUTSIDE; DROWNS OUT HUMAN VOICES, etc.

Some of these encyclopaedic features also apply straightforwardly to humans. Others (e.g. POWERFUL, GOES STRAIGHT AHEAD REGARDLESS OF OBSTACLES, PUSHES ASIDE OBSTRUCTIONS) have both a basic, physical sense and a further, psychologically-applicable sense, which may be frequently encountered and therefore often lexicalised. On the model of our discussion of 'cold', 'hard', 'rigid', etc. above, we would suggest that these words provide two potential inferential routes to the derivation of emergent properties. One possibility is that 'powerful', 'obstacle', etc. have both a basic physical sense and a broader, lexicalised, superordinate sense (POWERFUL*, OBSTACLE*, etc.) whose denotation includes both physical and psychological instances. The other is that 'powerful', 'obstacle', etc. have non-overlapping lexicalised physical and psychological senses (POWERFUL and POWERFUL**, OBSTACLE and OBSTACLE**, etc.), and that the interpretation of (9a) involves construction of a superordinate ad hoc concept (POWERFUL*, OBSTACLE*, etc.), which has both physical and psychological instances. In both cases, comprehension of (9a) would involve the use of BULLDOZER to derive contextual implications containing superordinate concepts such as POWERFUL*, OBSTACLE*, (etc.), which apply not only to bulldozers but to humans. As a result of the mutual adjustment process, Mary might therefore be understood as asserting that Robert is a BULLDOZER* (where a BULLDOZER*is POWER-FUL*, CRUSHING*, GOES AHEAD REGARDLESS OF OBSTACLES*, etc.) and implicating that she is reluctant to discuss her ideas with him because, being a BULLDOZER*, he is POWERFUL*, CRUSHING*, GOES AHEAD REGARDLESS OF OBSTACLES*, and therefore incapable of entering into a constructive, or mutually satisfactory, discussion. As with our previous analyses, both of these accounts are genuinely inferential: given the presumption of relevance conveyed by all utterances, interpretations along these lines are justified by the fact that Mary has produced this particular utterance in this discourse context.

5. CONCLUDING REMARKS

In the last section, we have given relevance-theoretic analyses of several metaphorical examples which have been widely seen in the literature as raising, in a particularly striking way, the issue of how hearers are able to recover emergent features of meaning. All other accounts we know of (specifically those in the cognitive linguistics literature) appeal to (non-inferential) associative mechanisms based on domain mappings of one sort or another (conceptual metaphors, blending of features from distinct domains). As noted above, on our inferential account, such associative links may arise as a result of repeated use of metaphors that bring together information from two distinct domains, and may in turn affect the outcome of the mutual adjustment process by altering the accessibility of contextual assumptions and implications; however, the resulting overall interpretation will only be accepted as the speaker's intended meaning if it satisfies the hearer's expectations of relevance and is properly warranted by the inferential comprehension heuristic discussed in section 2. As expected on a continuity account, other kinds of loose and even literal use can also give rise to emergent properties, although these are often less striking than in metaphorical examples such as (9)-(10). We claim that all these types of example are interpreted

by use of the same inferential comprehension procedure, with ad hoc concepts being added to the explicit content in the course of the mutual adjustment process so as to warrant the derivation of the expected contextual implications (or other cognitive effects).

We would like to suggest that the approach outlined here may shed interesting new light on the widespread phenomenon of polysemy in natural language (i.e. the fact that many or most linguistic forms have a range of distinct, though related, lexicalised senses). Double function adjectives such as 'cold', 'sharp', 'hard', and 'bright', some of which we discussed briefly in the last section, are a case in point, since both the psychological and the physical senses seem to be quite well-entrenched in the language and are likely to be lexicalised in the vocabulary of most speakers. Polysemy has been widely explored in frameworks that rely on systems of pre-established (non-inferential, associative) mappings between the elements of distinct cognitive domains (e.g. the physical and the psychological). We have suggested two possible inferential routes by which polysemy may arise.

In some cases, polysemy may arise through an inferential process of concept broadening, with the derived sense (e.g. COLD*, HARD*) being superordinate to the basic sense (e.g. COLD, HARD). In others, this superordinate sense may undergo a further inferential process of concept narrowing, yielding a distinct, non-overlapping basic sense (e.g. COLD**, HARD**) which may itself become lexicalised over time. In this way, inferential pragmatic processes of lexical narrowing and broadening may give rise to a range of related superordinate or non-overlapping lexicalised senses, with the appropriate analysis being established by empirical investigation on a case-by-case basis. This approach, based on a distinction between lexically encoded meanings and inferentially derived meanings which may in turn give rise to further encoded lexical meanings, might provide a useful theoretical framework for analysing not only polysemy but also lexical semantic change.

Notes

Much of the research for this paper was carried out as part of an AHRC-funded project 'A Unified Theory of Lexical Pragmatics' (AR16356). We would like to thank the Centre for the Study of Mind in Nature, University of Oslo, for support during the later stages of the research. We are grateful to our research assistants, Patricia Kolaiti, Tim

Wharton and, in particular, Rosa Vega Moreno, whose PhD work on metaphor we draw on in this paper. We are also grateful to Dan Sperber for many valuable discussions which have greatly enriched our views on metaphor, and to Elizabeth Camp and Sandra Lapointe for their editorial assistance and flexibility. An early version of this paper was published in 2006 in *Mind and Language* 21(3), 404-433; we have made significant revisions for this republication.

¹ The earliest treatment of metaphor within relevance theory (Sperber and Wilson 1985/6; Sperber and Wilson 1986/95) took a similar line, while differing substantially in its account of how the implicatures arose.

² Kolaiti (2005) reports several metaphorical uses of 'boiling' in the Bank of English corpus, describing e.g. rough seas, swirling clouds and chaotic thoughts. For a summary, see 'Corpus Analysis: An Overview' by the AHRC Lexical Pragmatics project team, accessible from the project home page at www.phon.ucl.ac.uk/home/lexprag07/corpus.html

³ Kolaiti (2005) shows that narrowing and broadening interact in quite complex ways in examples of this type. For instance, 'empty property' is sometimes used to mean a property empty of tenants (in which case it must be literally understood: even a single tenant is too many), and at other times to mean a property empty of furniture (in which case it may be understood as an approximation, since the odd piece of remaining furniture is inconsequential). See footnote 2 for details.

⁴ To save space, we present these simply as features rather than as complete propositions. However, since the function of encyclopaedic information is to provide premises for the derivation of contextual implications, each feature should be seen as a constituent of a complete proposition.

⁵ Here, the 'etc.' is intended to cover encyclopaedic features of strictly BOILING water that do not hold for broader interpretations; in (8c), it covers encyclopaedic features that hold both for strictly BOILING water and for water that is almost BOILING (i.e. BOILING*), but not for water that is BOILING** or BOILING***; and so on for (8b) and (8a). (See Sperber and Wilson 1998; Wilson and Sperber 2002; Wilson and Carston 2007.) We are not claiming, of course, that encyclopaedic information is neatly organised in this way: merely that the choice of a particular set of assumptions in the course of the mutual adjustment process will determine whether the utterance is intuitively 'literal', 'approximate', 'metaphorical', and so on.

⁶ As pointed out to us by Catherine Wearing, the word 'boiling' as understood in science does not necessarily entail the presence of heat (in conditions of sufficiently low pressure, water boils when it is cold). This is interesting, but has no particular consequences for our account of how (7) is ordinarily understood, although it does suggest that for some people, in some circumstances (rather few, we suspect), the literal, approximate and hyperbolic uses of 'boiling' would, like metaphorical uses, carry no implications of heat.

⁷ A further potentially important difference which we will not discuss here has to do with the nature of the superordinate categories themselves. The procedure that Glucksberg et al. propose for identifying appropriate superordinate categories (i.e. choose the first superordinate category that includes both the metaphor topic and the metaphor vehicle) suggests that they are treating metaphor as a type of category extension in which the superordinate category is arrived at by broadening alone (so in their example 'My job is a jail', the superordinate category JAIL* would include both the category of actual jails and the speaker's job). In the relevance-theoretic account, by contrast, superordinate categories may be produced either by broadening alone or by a combination of broadening and narrowing (so PRINCESS* in (1) includes some, but not all, actual princesses), and Glucksberg et al's procedure for finding an appropriate superordinate category no longer works straightforwardly.

⁸ These are fairly standardised cases, which might be listed in dictionaries as extra figurative senses. In dialects where 'bulldozer' has acquired an extra lexicalised sense (BULLDOZER*), the interpretation of (9a) would simply involve disambiguation plus contextual fine-tuning (see Vega Moreno 2004, 2007). From a historical point of view, however, this extra sense is likely to have become lexicalised as a result of repeated metaphorical uses of 'bulldozer' to convey the ad hoc concept BULLDOZER*, and it is this non-lexicalised type of case we are interested in analysing here. For expository purposes, it is convenient to use fairly standard examples that can be understood with a minimum of scene setting. In the case of novel metaphors, of course, the disambiguation account does not apply, and the only possible account is a wholly pragmatic one (for analysis of some novel uses, see Rubio Fernandez 2007; Sperber and Wilson 2008).

⁹ The same metaphor may, of course, receive quite different interpretations in other circumstances: it might be understood as suggesting, for instance, that men 'prey' on women in a way quite different from the preying of wolves on other creatures, or that humans will treat each other ruthlessly and mercilessly in extreme situations, and so on. The relevance-theoretic account sheds some light on the fact that even a conventional metaphor such as this one may be interpreted differently across hearers and times, since the accessibility of contextual assumptions varies across individuals and times.

¹⁰ For an interesting proposal to account for emergent properties by augmenting the relevance-theoretic account with the machinery of domain mappings, see Gibbs and Tendahl (2006).

¹¹ As shown by their stress patterns, the examples in (11a) and (11b) are productive adjective-noun combinations, whose semantic interpretations are derivable by compositional rules. By contrast, the semantic interpretations of compound nouns such as 'bird house' or 'house bird' are not systematically derivable, and we do not discuss them here.

¹² Sensory and kinaesthetic representations are themselves a source of conceptual information which may provide premises for inference.

¹³ For more detailed analysis of this example, see Vega Moreno (2004, 2007); Sperber and Wilson (2008). For an analysis using a cognitive linguistics framework, see Coulson and Oakley (2005).

¹⁴ In the framework of relevance theory, these are treated as *weak* implicatures, in the sense that no single one of them is essential to the interpretation and different hearers are likely to settle on slightly different ones. (On the derivation of weak implicatures, see Sperber and Wilson 2008.)

¹⁵ According to Tourangeau and Rips, the interpretation of (10) might involve a further narrowing, from the general concept VICIOUS (etc.), which applies to both wolves and (some) humans, to a more specific concept VICIOUS^{**} (etc.), which means VICIOUS IN THE HUMAN WAY (etc.). If so, we would treat it as resulting from inferential interaction between the metaphor topic ('men') and the constructed ad hoc concept WOLF^{*}, whose denotation would not itself be affected by the interaction.

¹⁶ For interesting discussion, which brings out the significance of Asch's work, see Rakova (2003).

 17 As this discussion shows, the interpretation of even a fairly standard metaphor such as 'Sally is a block of ice' is to some extent vague and open-ended, a point which is

35

often overlooked in philosophical and pragmatic accounts. For discussion of how this open-endedness can be dealt with in an inferential account, see e.g. Sperber and Wilson (1986/95: chapter 4, section 6); Pilkington (2000); Carston (2002a); Sperber and Wilson (2008).

¹⁸ It's worth noting that the Bank of English corpus (which contains 56 million words of naturally occurring text and discourse) contains only three metaphorical uses of 'bull-dozer': one a reference to a football player pushing people aside, and two references to Jacques Chirac being nicknamed 'the bulldozer' (Kolaiti 2005).

References

- Asch, Solomon E. 1955. On the use of metaphor in the description of persons. In Heinz Werner (ed.), *Expressive Language*, pp. 29–38. Worcester: Clark University Press.
- Asch, Solomon E. 1958. The metaphor: A psychological inquiry. In Renato Tagiuri and Luigi Petrullo (eds.), *Person Perception and Interpersonal Behavior*, pp. 86–94. Stanford: Stanford University Press.
- Barsalou, Lawrence. 1991. Deriving categories to achieve goals. In Gordon H. Bower (ed.), *The Psychology of Learning and Motivation*. New York: Harcourt Brace Jovanovich.
- Becker, Angela H. 1997. Emergent and common features influence metaphor interpretation. *Metaphor and Symbolic Activity* 12:243–259.
- Black, Max. 1962. Metaphor. In Max Black (ed.), Models and Metaphors, pp. 25–47. Ithaca, NY: Cornell University Press.
- Camp, Elisabeth. 2006. Contextualism, metaphor, and what is said. *Mind and Language* 21(3):280–309.
- Carston, Robyn. 1997. Enrichment and loosening: Complementary processes in deriving the proposition expressed? *Linguistische Berichte* 8:103–127.
- Carston, Robyn. 2002a. Thoughts and Utterances: The Pragmatics of Explicit Communication. Oxford: Blackwell.
- Carston, Robyn. 2002b. Linguistic meaning, communicated meaning and cognitive pragmatics. *Mind and Language* 17:127–48.
- Carston, Robyn. 2007. How many pragmatic systems are there? In Frapolli (2007), pp. 18–48.
- Coulson, Seana and Todd Oakley. 2005. Blending and coded meaning: Literal and figurative meaning in cognitive semantics. *Journal of Pragmatics* 37:1510–1536.
- Davis, Steven (ed.). 1991. Pragmatics: A Reader. Oxford: Oxford University Press.
- Fauconnier, Gilles and Mark Turner. 1998. Conceptual integration networks. *Cognitive Science* 22:133–187.
- Fauconnier, Gilles and Mark Turner. 2002. The Way We Think: Conceptual Blending and the Mind's Hidden Complexities. New York: Basic Books.
- Frapolli, Maria J. (ed.). 2007. Saying, Meaning, Referring: Essays on Francois Recanati's Philosophy of Language. Basingstoke: Palgrave Macmillan.
- Gernsbacher, Morton A., Boaz Keysar, Rachel Robertson, and Necia Werner. 2001. The role of suppression and enhancement in understanding metaphors. *Journal of Memory and Language* 45:433–450.

- Gibbs, Raymond W. 1994. The Poetics of Mind: Figurative Thought, Language and Understanding. Cambridge: Cambridge University Press.
- Gibbs, Raymond W. 1996. Why many concepts are metaphorical. Cognition 61:309-319.
- Gibbs, Raymond W. and Markus Tendahl. 2006. Cognitive effort and effects in metaphor comprehension: Relevance theory and psycholinguistics. *Mind and Language* 21(3):379–403.
- Gineste, Marie-Dominique, Bipin Indurkhya, and Veronique Scart. 2000. Emergence of features in metaphor comprehension. *Metaphor and Symbol* 15:117–135.
- Glucksberg, Sam. 2001. Understanding Figurative Language: From Metaphors to Idioms. Oxford: Oxford University Press.
- Glucksberg, Sam. 2004. On the automaticity of pragmatic processes: A modular proposal. In Noveck and Sperber (2004), pp. 72–93.
- Glucksberg, Sam and Zachary Estes. 2000. Feature accessibility in conceptual combination: Effects of context-induced relevance. Psychonomic Bulletin and Review 7:510–515.
- Glucksberg, Sam and Boaz Keysar. 1990. Understanding metaphorical comparisons: Beyond similarity. *Psychological Review* 97:3–18.
- Glucksberg, Sam, Deanna A. Manfredi, and Matthew S. McGlone. 1997. Metaphor comprehension: How metaphors create new categories. In Ward *et al.* (1997), pp. 327–350.
- Glucksberg, Sam, Mary Newsome, and Yevgeniya Goldvarg. 2001. Inhibition of the literal: Filtering metaphor-irrelevant information during metaphor comprehension. *Metaphor* and Symbol 15:277–293.
- Grice, H. Paul. 1967/89. Logic and conversation. In H. Paul Grice (ed.), Studies in the Way of Words. Cambridge, MA: Harvard University Press. William James Lectures 1967.
- Hampton, James. 1997. Emergent attributes in combined concepts. In Ward *et al.* (1997), pp. 83–110.
- van der Henst, Jean-Baptiste and Dan Sperber. 2004. Testing the cognitive and communicative principles of relevance. In Noveck and Sperber (2004), pp. 141–171.
- Kintsch, Walter. 2000. Metaphor comprehension: A computational theory. *Psychonomic Bulletin and Review* 7:257–266.
- Kolaiti, Patricia. 2005. 'The empty vessel makes the greatest sound': Corpus analysis and lexical pragmatics. Paper delivered to the AHRC Conference on Word Meaning, Concepts and Communication, Cumberland Lodge, Windsor Great Park, September 2005.
- Lakoff, George. 1987. Women, Fire and Dangerous Things. Chicago: Chicago Univ. Press.
- Lakoff, George. 1994. Conceptual Metaphor Home Page. Available at: http://cogsci.berkeley.edu/lakoff/MetaphorHome.html.
- Lewis, David. 1975. Languages and language. Language, Mind and Knowledge. Minnesota Studies in the Philosophy of Language 7:3–35. Reprinted in Lewis 1983, 163-188.
- Lewis, David. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8:339–359. Reprinted in Lewis 1983, 233-249.
- Lewis, David. 1983. Philosophical Papers, volume 1. Oxford: Oxford University Press.
- Martinich, Al P. 1984/91. A theory for metaphor. *Journal of Literary Semantics* 13:35–56. Reprinted in Davis (ed.) 1991, 507-518.
- McGlone, Matthew. S. and Deanna A. Manfredi. 2002. Topic-vehicle interaction in comprehension. *Memory and Cognition* 29:1209–1219.

Noveck, Ira, Maryse Bianco, and Alain Castry. 2001. The costs and benefits of metaphor.

Metaphor and Symbol 16:109–121.

- Noveck, Ira and Dan Sperber (eds.). 2004. *Experimental Pragmatics*. Basingstoke: Palgrave Macmillan.
- Pilkington, Adrian. 2000. Poetic Effects: A Relevance theory Perspective. Amsterdam: John Benjamins.
- Pugmire, David. 1998. Rediscovering Emotion. Edinburgh University Press.
- Rakova, Marina. 2003. The Extent of the Literal: Metaphor, Polysemy and Theories of Concepts. Basingstoke: Palgrave Macmillan.
- Recanati, Francois. 1995. The alleged priority of literal interpretation. *Cognitive Science* 19:207–232.
- Recanati, Francois. 2002. Does linguistic communication rest on inference? *Mind and Language* 17:105–126.
- Recanati, Francois. 2004. Literal Meaning. Cambridge: Cambridge University Press.
- Rips, Lance. 1995. The current status of research on concept combination. *Mind and Language* 10:72–104.
- Romero, Esther and Belen Soria. 2007. A view of novel metaphor in the light of Recanati's proposals. In Frapolli (2007), pp. 145–159.
- Rubio Fernandez, Paula. 2005. Pragmatic Processes and Cognitive Mechanisms in Lexical Interpretation: The On-Line Construction of Concepts. Ph.D. thesis, University of Cambridge.
- Rubio Fernandez, Paula. 2007. Suppression in metaphor interpretation: Differences between meaning selection and meaning construction. *Journal of Semantics* 24(4):345– 372. Special Issue on Processing Meaning.
- Sperber, Dan and Deirdre Wilson. 1985/6. Loose talk. Proceedings of the Aristotelian Society 86:153–171. Reprinted in Davis (ed.) 1991, 540-549.
- Sperber, Dan and Deirdre Wilson. 1986/95. *Relevance: Communication and Cognition*. Oxford: Blackwell. Second edition 1995.
- Sperber, Dan and Deirdre Wilson. 1998. The mapping between the mental and the public lexicon. In Peter Carruthers and Jill Boucher (eds.), *Language and Thought*, pp. 184– 200. Cambridge: Cambridge University Press.
- Sperber, Dan and Deirdre Wilson. 2008. A deflationary account of metaphors. In Raymond W. Gibbs (ed.), *Cambridge Handbook of Metaphor and Thought*, pp. 84–105. New York: Cambridge University Press.
- Tourangeau, Roger and Lance Rips. 1991. Interpreting and evaluating metaphors. *Journal* of Memory and Language 30:452–472.
- Vega Moreno, Rosa. 2004. Metaphor interpretation and emergence. UCL Working Papers in Linguistics 16:297–322.
- Vega Moreno, Rosa. 2007. Creativity and Convention: The Pragmatics of Everyday Figurative Speech. Amsterdam: John Benjamins.
- Ward, Thomas, Steven Smith, and Jyotsna Vaid (eds.). 1997. Creative Thought: An Investigation of Conceptual Structures and Processes. Washington DC: American Psychological Association.
- Wilson, Deirdre. 2003. Relevance theory and lexical pragmatics. Italian Journal of Linguistics / Rivista di Linguistica 15:273–291. (Special Issue on Pragmatics and the Lexicon).
- Wilson, Deirdre and Robyn Carston. 2007. A unitary approach to lexical pragmatics: Relevance, inference and ad hoc concepts. In Noel Burton-Roberts (ed.), *Pragmatics*, pp. 230–259. London: Palgrave.

Wilson, Deirdre and Dan Sperber. 2002. Truthfulness and relevance. *Mind* 111:583–632.
Wilson, Deirdre and Dan Sperber. 2004. Relevance theory. In Laurence Horn and Gregory Ward (eds.), *Handbook of Pragmatics*, pp. 607–632. Oxford: Blackwell.