Suppression in Metaphor Interpretation: Differences between Meaning Selection and Meaning Construction

PAULA RUBIO FERNÁNDEZ University College London

Abstract

Various accounts of metaphor interpretation propose that it involves constructing an ad hoc concept on the basis of the concept encoded by the metaphor vehicle (i.e. the expression used for conveying the metaphor). This paper discusses some of the differences between these theories and investigates their main empirical prediction: that metaphor interpretation involves enhancing properties of the metaphor vehicle that are relevant for interpretation, while suppressing those that are irrelevant. This hypothesis was tested in a cross-modal lexical priming study adapted from early studies on lexical ambiguity. The different patterns of suppression of irrelevant meanings observed in disambiguation studies and in the experiment on metaphor reported here are discussed in terms of differences between meaning selection and meaning construction.

1 CONCEPT CONSTRUCTION IN METAPHOR INTERPRETATION

One of the current directions in metaphor theory is based on the assumption that metaphor comprehension involves both the use of previously acquired concepts and schemas and the creation of new ones (Cacciari & Glucksberg 1994). Glucksberg & Keysar (1990) were among the first researchers to develop this idea with their classinclusion model of metaphor interpretation, moving away from the traditional Aristotelian view that nominal metaphors of the form 'X is a Y' are understood as implicit comparisons (i.e. 'X is like a Y') (Bowdle & Gentner 2005). In their view, nominal metaphors are 'exactly what they appear to be: class-inclusion assertions' (Glucksberg & Keysar 1990: 3).

According to the class-inclusion model, the content of a metaphor is assigned to a 'diagnostic category' that is labelled by the figurative expression or 'metaphor vehicle', whose literal denotation is also a stereotypical subset of that category. Consider the following example from Glucksberg and Keysar (1990):

(1) My job is a jail.

Such a metaphorical statement would be a true class-inclusion assertion, where 'jail' not only has its usual denotation but also refers to a broader ad hoc category including both jails and the speaker's job. Glucksberg *et al.* (1997a) refer to this aspect of metaphor interpretation as 'dual reference'. Thus, 'jail', as the metaphor vehicle, would give a name to the new category while its salient properties would become prototypical of that category. Since 'my job', as the topic of the metaphor, is asserted to be a member of the ad hoc category, typical properties of jails, such as 'being unpleasant', 'constraining and punishing', would now be attributed to it.

Given their experimental approach, Glucksberg and his colleagues make precise empirical predictions about the interpretation and processing of metaphors. For example, interpreting a nominal metaphor like (1) involves enhancing attributes of the vehicle 'jail' that are appropriate for the topic 'my job' (e.g. 'unpleasant', 'constraining') while suppressing those attributes that are inappropriate (e.g. 'has bars on the windows'). This hypothesis has been tested in a number of experiments (see below for review).

Another account of metaphor interpretation based on ad hoc concept construction is that of Relevance Theory (Carston 1996, 2002; Sperber & Wilson 1997, 2002; for an earlier relevance-theoretic account, entirely in terms of implicature, see Sperber & Wilson 1986/ 1995). In this theoretical framework, metaphor is seen as a type of 'loose use of language' comparable to various other phenomena usually discussed in other terms (e.g. approximations, category extensions, hyperboles and neologisms; see Wilson 2003 for examples and discussion). This unified account of loose use falls within the relevance-theoretic programme for 'lexical pragmatics', understood as the branch of linguistics that investigates the pragmatic processes by which the meaning communicated by the use of a word in context comes to differ from the linguistically encoded meaning of that word.

Relevance Theory distinguishes two cases where the concept communicated by the use of a word differs from the concept encoded by that word: concept narrowing and concept loosening, both resulting from a general pragmatic process of 'concept adjustment' (Carston 2002; Wilson 2003). In instances of narrowing, a word is used to convey a more specific concept than the one it encodes, whereas in instances of 'loosening', the concept communicated is more general than the lexical concept. The interpretation of most nominal metaphors involves a combination of narrowing and broadening (see Carston 1996, 2002 for a discussion of the various types of ad hoc concept constructed in metaphor interpretation). Consider the following example (adapted from Carston 2002):

(2) Caroline is a princess.

Suppose that (2) was used to convey that Caroline is spoilt. In interpreting this nominal metaphor, an ad hoc concept PRINCESS* would be constructed, resulting from a process of loosening and narrowing the lexical concept PRINCESS. On the one hand, Caroline would be included in the extension of PRINCESS*, which means that the property 'female royal of a certain sort', which is definitional of the lexical concept PRINCESS, would not characterize all members of the new category given that Caroline (and other spoilt women included in the ad hoc category) is not a royal. On the other hand, the entities falling within the denotation of PRINCESS* might be characterized as having the property spoilt. Those princesses who do not behave in this way would not be included in the extension of the ad hoc concept. Therefore, in the process of fine-tuning the interpretation of the word 'princess', the extension of the communicated concept PRINCESS* would have been not only broadened but also narrowed in relation to that of the encoded concept.

Recanati (2004) offers an account of metaphor interpretation that is similar to Carston's. According to Recanati, a concept has certain 'conditions of application' associated with it. Concept narrowing would consist in restricting the reference of a concept by contextually providing further conditions that are not linguistically encoded. Concept loosening is understood as the converse of the process of narrowing: some condition of application packed into the concept literally expressed by a word is contextually dropped so that the extension of the concept is widened (Recanati 2004). Despite the clear similarities in their accounts, there are some major differences between Recanati's and the relevance-theoretic view of metaphor interpretation, which will be discussed in section 2.

Although both Relevance Theory and Recanati's framework are open to empirical investigation (Recanati 2004; Wilson & Sperber 2004), they do not make specific empirical predictions about metaphor processing and interpretation. However, Carston (2002) and Recanati (2004) describe narrowing and loosening in terms of property promotion and demotion, which could be understood in terms of degree of activation of the properties in question (Rubio 2005). According to these authors, if an ad hoc concept results from a process of narrowing, some encyclopaedic property of the lexical concept is elevated to the status of a logical or content-constitutive component. In example (2) above, 'princess' would express PRINCESS*, an ad hoc concept resulting from strengthening the lexical concept PRINCESS by making the encyclopaedic property 'spoilt' content constitutive of the new concept (i.e. a princess who is not spoilt would not fall under the ad hoc category PRINCESS*). On the other hand, if an ad hoc concept results from loosening a lexical concept, one or more of the logical or definitional properties of the lexical concept are discarded (Carston 1996, 2002; Recanati 2004). In example (2), the contentconstitutive property 'female royal of a certain sort' would be dropped from the logical entry of the lexical concept PRINCESS in order to include Caroline, who is not a royal, under the ad hoc concept PRINCESS*.

Similarly to the proposal made by Glucksberg and his colleagues, it is natural to suppose that those encyclopaedic properties of the lexical concept which are promoted to the status of content constitutive of the resulting ad hoc concept become and remain active during the metaphor interpretation process given their contextual relevance (Rubio 2005). Conversely, the activation of those logical properties that have been demoted in concept loosening may be suppressed during processing (Recanati 2004; Rubio 2005). The study reported in this paper was aimed at testing this twofold empirical prediction that is common to the class-inclusion model, the relevance-theoretic account and Recanati's view of metaphor interpretation.

2 FURTHER THEORETICAL ISSUES

Although the models of metaphor interpretation discussed above share the view that metaphor comprehension involves enhancing relevant properties of the vehicle while suppressing irrelevant ones, their accounts are different in some fundamental respects. In particular, there are important differences between Recanati's and the relevancetheoretic views of metaphor interpretation (for a discussion of the differences between the class-inclusion and the Relevance Theory models of metaphor interpretation, see Rubio 2005). Although these differences might not be open to direct empirical investigation, the results of the present study might shed some light on some of these issues.

According to Recanati (2004), narrowing and loosening are 'primary pragmatic processes'. This type of pragmatic process is characterized by being prepropositional (i.e. they do not require that a proposition has been identified already), unconscious (i.e. normal interpreters are not aware of their local operation) and purely associative (i.e. they operate in a blind mechanical fashion that involves no inferential process on the part of the hearer) (Recanati 2004). Relevance Theory does not make a distinction between primary and secondary pragmatic processes, advocating a unified account of the pragmatic processes involved in interpretation (Carston forthcoming). In particular, according to Relevance Theory, all pragmatic processes are uniformly inferential rather than some being merely associative (i.e. the input and output of pragmatic computations are related as premise and conclusion in an argument), and all are constrained by considerations of relevance (i.e. they tend to maximize cognitive effects while minimizing processing effort).

Unlike the relevance-theoretic account, Recanati's model of lexical interpretation is based solely on accessibility: in processing a word, its literal interpretation is accessed first and triggers the activation of associatively related representations (Recanati 2004). Both the literal concept activated by the linguistic expression and some of the other representations activated by association are possible candidates for the concept that will be selected for interpretation. Although these associated candidates are generated via, hence after, the literal concept, they are all processed in parallel and compete for activation. The representation that is most active or accessible when the interpretation process stabilizes will be selected and undergo semantic composition with the other components of the utterance, while all other candidates for the meaning of the word are suppressed (Recanati 1995, 2004).

Consider the following example of metonymy (discussed in Recanati 1995, 2004):

(3) The ham sandwich has left without paying.

In Recanati's view, the expression 'the ham sandwich' first receives its literal interpretation by activating the representation of HAM SANDWICH. Then activation spreads from the literal representation to other associated representations, such as that of HAM SANDWICH ORDERER. 'All these representations activated by the description "the ham sandwich" contribute potential candidates for (...) going into the interpretation of the global utterance' (Recanati 2004: 29). Although the literal representation of HAM SANDWICH might initially have been

a more accessible candidate than HAM SANDWICH ORDERER receives increased activation further down the line because of its suitability as an argument of the predicate 'leave without paying', so it ends up being the most accessible of the candidates once the entire utterance has been processed. Therefore, the non-literal candidate that derived from the literal one will be retained as part of the proposition expressed by (3), and the literal interpretation is suppressed (Recanati 2004).

Recanati's model is based on the notion of 'accessibility shifts': although a given representation of the meaning of some linguistic expression may be the most accessible one at some point in processing (e.g. the literal interpretation at an early stage in processing), activation levels may change and another representation (e.g. the derived nonliteral interpretation) may become the most accessible one at some later point in processing. Recanati (1995, 2004) seems to distinguish two possible factors that might determine these accessibility shifts. The first one is the processing of further linguistic material, as in the example just discussed, where the initial accessibility ordering is reversed once the predicate expression has been processed. The second factor is world knowledge structures or 'schemata': according to Recanati, interpretation is generally driven by schemata, so utterance interpretation is to a large extent a top-down process driven by world knowledge (Recanati 2004). In contrast, in the relevance-theoretic account, speaker's intentions and not only world knowledge structures are among the factors that can affect the accessibility of a word's interpretation (Sperber & Wilson 2002).

Finally, according to Recanati (2004), a global literal interpretation does not necessarily precede figurative interpretation (cf. Grice 1975, 1989). However, at the local lexical level, the literal interpretation of the constituents of a metaphorical expression is accessed before the figurative interpretation of the utterance is derived (cf. Sperber & Wilson 2002). Nevertheless, Recanati (2001) considers a case where the literal interpretation of a figurative expression may not necessarily be suppressed once the intended interpretation of the utterance has been accessed: when a certain threshold of awareness is reached in interpreting an expression metaphorically, it results from some dimensions of the literal meaning remaining active even if they have been filtered out of the interpretation would still need to be the most accessible one in order to be selected.

Although these two theoretical issues are probably not empirically testable in any direct way, the experimental study reported below will give some hints as to whether a figurative lexical interpretation is derived associatively or inferentially and whether it is necessarily the most accessible interpretation once it is selected.

3 METAPHOR PROCESSING AND LEXICAL AMBIGUITY

The predictions of the class-inclusion model of metaphor interpretation were tested in a series of experiments (e.g. Gernsbacher *et al.* 1995; Glucksberg *et al.* 1997b, 2001; Gernsbacher *et al.* 2001). The initial hypothesis was that interpreting a metaphor (e.g. 'That defence lawyer is a shark') involves enhancing attributes of the vehicle that are appropriate for the topic ('aggressive', 'vicious', 'tenacious') while suppressing those attributes that are inappropriate ('swims', 'has fins', 'lives in the ocean'). Glucksberg, Gernsbacher and their colleagues used a sentence verification task, where subjects had to verify a metaphorical statement and then evaluate a second assertion that could be related to either the appropriate aspects of the metaphor vehicle ('Sharks are tenacious') or the inappropriate ones ('Sharks are good swimmers').

The pattern of results observed in these studies seemed to confirm the predictions made by Glucksberg and colleagues: after reading a metaphorical statement, subjects were faster to verify statements related to the metaphoric interpretation of the vehicle and slower to verify statements related to the literal interpretation of the vehicle, when compared to the time it took them to respond to the same statements after reading a control sentence. These data therefore suggest that interpreting a metaphor does involve enhancing attributes that are relevant to the metaphorical interpretation of the vehicle while suppressing those that are irrelevant.

However, the experimental design used in these studies makes it difficult to evaluate the power of their results, especially concerning the question whether suppression reduces the activation of metaphorirrelevant properties below baseline: the type of control sentences used in these experiments included the metaphor vehicle, which was the last word in the sentence in both the critical and the control conditions (e.g. 'That hammerhead is a shark'). The control sentences were therefore not properly unrelated to the target sentences ('Sharks are tenacious'/'Sharks are good swimmers') so the degree of activation of the target properties ('tenacious'/'swims') after processing the control sentences would not correspond to a 'zero' level of activation. Because suppression is defined with reference to the baseline condition, an underestimate of baseline performance would result in an overestimate of suppression. It follows that this type of baseline measure would not allow for making an accurate distinction between low activation and below-baseline suppression (see Rubio 2005 for further discussion).¹ Nonetheless, McGlone & Manfredi (2001) also obtained results supporting the predictions of the class-inclusion model in a similar study where they controlled for the repetition of the metaphor vehicle between prime and target sentences, but using a baseline condition that did not prime the relevant properties (see reference for details of the experimental design).

Overall, the results of previous studies investigating the predictions of the class-inclusion model suggest that both enhancement and suppression are indeed involved in metaphor interpretation, although the type of controls used makes it difficult to evaluate the power of such mechanisms. The control condition in the present study included completely unrelated targets. However, rather than the sentence verification task used by Glucksberg, Gernsbacher and their colleagues, I used a cross-modal lexical priming paradigm, which allows making an on-line measure of property activation across time. This in turn would allow investigating at which point in processing suppression dampens the activation of irrelevant literal properties of the metaphor vehicle. This question is related to the role of conscious, attentional processes in metaphor processing.

The priming paradigm used in this study was adapted from early studies of lexical ambiguity (Swinney 1979; Tanenhaus *et al.* 1979; Onifer & Swinney 1981). In the original experiments, participants were presented with sentences in the acoustic modality which included a homonym (e.g. 'The man found several bugs in his room'). At the offset of the ambiguous prime, participants had to make a lexical decision on a visual target. Critical targets could be related to either of the two meanings of the homonym (e.g. 'spy' or 'ant' in the above example). Facilitation relative to an unrelated control was interpreted in terms of meaning activation. The results of these experiments showed an early activation of target words related to both meanings of the

¹ In the third experiment reported in Gernsbacher *et al.* (2001), they used properly unrelated control sentences. However, in this experiment, rather than being slower, participants were actually faster at verifying a metaphor-irrelevant statement after reading a metaphor than after reading the unrelated control. Gernsbacher and her colleagues explain this facilitation as a result of the imbalance in the repetition effect between the critical and the control conditions. It is unclear, however, to what extent the standardized scores that had to be computed in order to observe suppression of metaphor-irrelevant information accounted exclusively for the uneven repetition of the metaphor vehicle across conditions and offered an accurate measure of suppression. Glucksberg *et al.* (2001) modified the original design of the study in order to avoid the metaphor vehicle and would therefore have prime the critical properties (see reference for details).

homonym, which was interpreted in terms of an automatic, exhaustive process of spreading activation of associates (Schvaneveldt & Meyer 1973). However, the activation of the contextually inappropriate meaning dropped as early as 200–300 ms from the offset of the ambiguous word. This pattern of results was interpreted as showing active suppression of the irrelevant reading of the ambiguity, given that passive decay should take considerably longer (Neely 1976; Tanenhaus *et al.* 1979). However, since controlled, attentional processes take 400–500 ms to operate (Posner & Snyder 1975; Neill *et al.* 1995), these authors argue that although the meaning selection process must be context-sensitive (unlike the early spreading activation phase), it operates in an almost automatic way (unlike later conscious processes) (see Shiffrin & Schneider 1977). This would explain why hearers are usually unaware of having encountered a homonym in a disambiguating context (Gernsbacher 1990).

In the study reported in this paper, participants were presented with contexts biased in favour of metaphorical interpretations, for example 'Nobody wanted to run against John at school. John was a cheetah.' Critical targets for a lexical decision were either metaphor-inconsistent properties of the metaphor vehicle ('cat') or metaphor-relevant properties ('fast').² As in the above studies, I took facilitation relative to an unrelated control as indicative of property activation. In order to investigate the time course of activation, targets were presented 0, 400 and 1000 ms from the offset of the metaphoric prime. Assuming that metaphor-relevant properties are enhanced in metaphor interpretation, whereas metaphor-inconsistent properties are suppressed (Glucksberg & Keysar 1990; Carston 2002; Recanati 2004), the question that needs to be addressed is at which point in processing the activation level of these two types of associates diverges. In particular, it would be interesting to see whether metaphor-inconsistent properties would be suppressed as early as the irrelevant meanings of homonyms (i.e. 200-300 ms from the offset of the metaphor vehicle), or whether their suppression would involve later, more attentional processes on the part of the interpreter.

Although suppression would be involved in both disambiguation and metaphor interpretation (Gernsbacher & Faust 1991), these linguistic phenomena involve different pragmatic processes: meaning selection and meaning construction, respectively (see Recanati 1995;

² Although the superordinate 'cat' is not strictly speaking a property of CHEETAH, I am taking the activation of this associate as indicative of the activation of 'being a cat', which is a content-constitutive property of the prime concept.

Carston 2002). It is therefore possible that the mechanism of suppression operates differently in each case. Disambiguating a homonymous word like 'bank' would involve selecting one of its two meanings (i.e. financial institution or side of a river). Given that these two meanings are part of the mental lexicon of the hearer, in processing the lexical form 'bank' activation would spread to two different lexical entries (for the sake of simplicity, BANK-1 and BANK-2). In contrast, interpreting for the first time a metaphor like 'Mary is a nightingale' would involve constructing an ad hoc concept NIGHTINGALE* from the lexically encoded concept NIGHTINGALE (Carston 2002; Glucksberg 2003; Recanati 2004). It would therefore be interesting to see whether these differences in the accessibility of the various concepts involved in disambiguation and metaphor interpretation may have an effect on the operation of the mechanism of suppression.

To summarize, the following study was aimed at investigating two questions: first, the empirical prediction of various theories that metaphor interpretation involves enhancing metaphor-relevant properties of the vehicle while suppressing metaphor-irrelevant ones and second, at which point in processing suppression reduces the activation of metaphor-irrelevant information in relation to the suppression of contextually inappropriate meanings in early studies of lexical ambiguity. Also, some secondary theoretical issues about the accessibility of the figurative interpretation of the metaphor vehicle will be discussed.

4 AN ON-LINE STUDY OF METAPHOR PROCESSING

4.1 Method

4.1.1 *Participants* The participants in this experiment were 60 undergraduate students at Cambridge University and University College London who volunteered to take part in the experiment. They all had English as their first language. Each session lasted approximately 15 min.

4.1.2 *Materials and design* A set of 22 common nouns with predictable superordinates and distinctive properties were selected as primes. Two questionnaires based on the literature on prototypes (Rosch & Mervis 1975; Barsalou 1987) offered a direct account of property dominance. After piloting the questionnaires on 15 participants, the final version was distributed among 65 participants. Having chosen a list of words with predictable superordinates and

distinctive properties, the results were as expected apart from two terms, 'tip' and 'spring', which are ambiguous and did not elicit a uniform response. These two terms were discarded. For each of the 20 remaining concepts, the most frequent superordinate term was chosen from the brief definition task in the first questionnaire. Likewise, the most frequent distinctive property was selected from the three different tasks in the two questionnaires (i.e. a brief definition task, a distinctive property listing and a free-association task; for the set of primes and targets, see Appendix A). Both types of target were therefore strong associates of the prime concepts (what I have also called 'core features' of the primes; see Rubio 2005, forthcoming).

A metaphorically biased context was constructed for each one of the 20 primes so that superordinates were inconsistent with the figurative interpretation, whereas distinctive properties were relevant for interpretation. Each context ended in a nominal metaphor of the form 'X is a Y', Y being always the prime concept. Because the metaphors were novel, the preceding context included two sentences on average to make sure that the nominal metaphor would be comprehensible (for the set of metaphoric contexts, see Appendix B). The 20 metaphoric contexts were divided into two equal groups matched for word length and frequency of the corresponding superordinates and distinctive properties (Johansson & Hofland 1989). One group of contexts were paired with related superordinates and distinctive properties. For the other group, targets were scrambled so the sentences were paired with unrelated target words. The unrelated contexts and targets served as controls. Two lists of materials were constructed by pairing one group of contexts with related targets in List A and with unrelated targets in List B, and the other group of contexts with unrelated targets in List A and with related targets in List B. Another set of 20 metaphoric contexts was constructed and paired with English-like non words. Critical and filler sentences were randomized individually for each participant in each list of materials. Participants were randomly assigned to one Target Type, List and Inter Stimulus Interval (ISI), so each participant saw each context and the corresponding target only once.

Given that a word is identified at the point in time when acoustic information uniquely specifies it, which may actually occur before the physical ending of the word (Marslen-Wilson 1987), for each of the 20 nouns in our materials a point was selected where the prime would be unequivocally recognized. Targets were presented visually at the end of the acoustic signal 0, 400 or 1000 ms after the word-recognition point selected for each prime. This enabled accurate measuring of initial semantic activation, while controlling for the possibility that an early contextual effect may result from an early word recognition followed by a fast property selection given the length of the primes.

Sentences were recorded at a normal rate by a male speaker on an Apple Macintosh computer. The auditory stimuli and the visual targets were synchronized using a specialized computer program.

The experimental items were preceded by two sets of practise trials. The first one consisted of a lexical decision task and the second one included both sentential contexts in the acoustic modality and visual targets for a lexical decision. The latter contained six metaphoric contexts similar to the critical ones, although the corresponding visual targets were not related to the primes in any of the practice trials.

4.1.3 *Apparatus* The experiment was conducted on a Toshiba laptop computer. The sentences were presented through a pair of headphones plugged into the laptop. The visual probes were presented in capital letters in the middle of the computer screen on a white background. Responses to the visual targets were made via a response box connected to the laptop. 'Word' responses were made with the thumb or the index finger of the right hand and 'non-word' responses with the thumb or the index finger of the left hand. Target words remained on the screen until the participant had made a decision. There was a 1000-ms delay between the offset of the visual target and the onset of the following acoustic context.

4.1.4 *Procedure* The experiment was presented to the participants as a simple psycholinguistic experiment investigating the interpretation of metaphorical language. Participants were told that they would be listening to a series of short texts through the headphones and that each text would end in a metaphor. To make sure that participants derived the intended interpretation of the metaphor, they were asked to try to visualize the figurative meaning of the metaphor. At the end of each sentence, a string of letters would appear on the computer screen and they should try to indicate as fast and accurately as possible whether the string of letters was a word of English or not by pressing the corresponding key on the response box.

Participants were first given standard written instructions, which were then explained individually by the experimenter. It was emphasized that both tasks (i.e. listening carefully to the sentences and making a fast lexical decision) were equally important, although they should be taken as independent tasks.

Participants were tested individually. They ran through the two sets of practice trials with the experimenter and got appropriate feedback on their performance. When being tested on the critical materials, participants were left on their own in a closed room or cubicle.

In order to make sure that participants paid adequate attention to the contexts, a short memory test was given at the end of the experiment. Participants had been told about this memory test in the instructions.

4.2 Results

The minimum of correct responses required in the memory test was 2.5 standard deviations below the participants' average of correct responses per ISI. No participant had to be replaced for failing to meet this criterion.

The mean response time, standard deviation and proportions of missing data for each Relatedness condition, together with the facilitation (i.e. the difference between the experimental [related] and the control [unrelated] conditions) and its significance level per Target Type and ISI are presented in Table 1. A response time data point was treated as 'missing' if it was either from an erroneous response or over 2.5 standard deviations above the participant's average response time to the word targets in his exercise.

The activation curves of superordinates and distinctive properties in metaphoric contexts, understood as the priming effect observed across the three ISIs for each target type, are given in Figure 1.

The statistical analysis of the data examined the effects of Target Type (superordinate/distinctive property), Target Relatedness (related/ unrelated), ISI (0/400/1000 ms) and List (A/B). Mean reaction times were entered into four-way analyses of variance (ANOVAs), with participants (F_1) as the random variable.³ There was a significant main effect of Relatedness, $F_1(1,48) = 46.03$, MSE = 1102.1, P < 0.001; ISI, $F_1(2,48) = 4.425$, MSE = 59251, P < 0.02 and Target Type, $F_1(1,48) = 4.968$, MSE = 59251, P < 0.04. The fastest reaction times were observed in the related condition (42 ms difference), the 400-ms condition (93 ms difference). Only the 2 × 2 × 3 × 2 interaction (Relatedness × Target Type × ISI × List) was significant, $F_1(2,48) = 6.381$, MSE = 1102.1, P < 0.004. The effect of List was not systematic,

³ Because of the small scale of the experiment, the design was not powerful enough to carry out reliable analyses per item. However, List was included as an independent variable to see whether the distribution of the materials had had any significant effect on the ANOVAs. Nonetheless, without an item analysis it is not possible to establish whether the results generalize to other metaphors, which is obviously a limitation of the study.

2 - 0	· ·		X Z Z	· ·
55×	Viinneocoron	110	Matanhan	ntornrotation
.).)()	SUDDIESSION	111	wietablior	Interpretation
			r	

	Relatedness	ISI			
Target		0	400	1000	
Superordinates	Related Unrelated Facilitation	883 (198, 0.03) 919 (187, 0.05) -36*	644 (162, 0.03) 698 (163, 0.07) -55***	791 (167, 0.02) 799 (156, 0.03) -8	
Distinctive properties	Related	740 (243, 0.04)	658 (58, 0.04)	598 (104, 0.03)	
	Unrelated Facilitation	782 (230, 0.1) -42**	710 (79, 0.07) -52***	651 (123, 0.06) -53**	

 Table 1
 Mean reaction times (in milliseconds), standard deviations, proportions of missing data and facilitation for each condition of the experiment

*P < 0.1 **P < 0.05 ***P < 0.01.

given that the highest facilitation for each Target Type was not observed for the same List condition across ISIs. This interaction with List could therefore be related to the different speed of reaction of the different groups of participants tested in each condition.

A 2 × 2 × 3 × 2 (Relatedness × Target Type × ISI × List) ANOVA was carried out on the arcsine transformation of the missing data using participants as the random factor. The missing data were arcsine transformed to stabilize variances (Winer 1971). Only a significant main effect of Relatedness was observed, $F_1(1,48) = 7.031$, MSE = 0.029, P < 0.02. The highest missing rate was observed in the unrelated condition (0.179), so the facilitation observed could not have been due to missing data points lowering the average response time in the related condition. No interaction reached significance level (all $F_1 < 1.9$).

Given that the greatest difference in the level of priming for superordinates and distinctive properties was observed at the longest delay (45 ms difference), a 2 × 2 × 2 (Relatedness × Target Type × List) ANOVA was carried out on the reaction time data for the 1000-ms condition. This analysis is particularly relevant for the investigation of attentional processes, which take place around 500 ms from the offset of the prime (e.g. Neely 1976; Yee 1991). There was a significant main effect of Relatedness, $F_1(1,16) = 8.705$, MSE = 1078.9, P < 0.01, and Target Type, $F_1(1,16) = 7.004$, MSE = 41614, P < 0.02, with the fastest reaction times being observed in the related condition (31 ms difference) and the distinctive properties condition (171 ms difference). The Relatedness × Target Type interaction, which



Figure 1 Activation curves of superordinates and distinctive properties in metaphoric contexts.

was critical for the present investigation, was significant, $F_1(1,16) = 4.559$, MSE = 1078.9, P < 0.05, with the highest level of priming being observed in the distinctive property condition. The Relatedness × Target Type × List interaction was also significant, $F_1(1,16) = 11.507$, MSE = 1078.9, P < 0.005. The highest level of priming was observed on the List A condition for superordinates (-30 ms) and on the List B condition for distinctive properties (-101 ms). Since this pattern of results was not consistent across ISIs in previous analyses and List did not show a significant main effect, this interaction with List could be related to the different performance of the different groups tested in each condition rather than to the particular distribution of the materials. The corresponding ANOVA of the arcsine transformation of the missing data did not show any significant results (all $F_1 < 1.5$).

4.3 Discussion

Despite their different roles in metaphor interpretation, both superordinates and distinctive properties were active up to 400 ms from the offset of the metaphor vehicle, their levels of activation deviating only at the 1000-ms delay. However, according to the results of the ANOVAs, the activation patterns of these associates were different, especially at the longest delay, where the critical Relatedness \times Target Type interaction was significant.

It looks as if the loss of activation of superordinates between 400 and 1000 ms must be the result of active suppression of metaphorinconsistent information. In a previous study, we observed that superordinates remain active up to 1000 ms in literal neutral contexts, where no property of the prime was particularly relevant or irrelevant for interpretation (Rubio *et al.* 2003). Their loss of activation in metaphoric contexts could not, therefore, have been the result of passive decay but has to be due to active suppression.

The suppression of superordinates between 400 and 1000 ms in metaphoric contexts would have been due to attentional processes different from the processes involved in lexical disambiguation, which occur as early as 200 ms from the offset of the homonym. The main differences are that attentional processes (i) involve some level of awareness, (ii) need time to develop and (iii) are of limited capacity (Keele & Neill 1978; Neill & Westberry 1987). Therefore, although the results support the theoretical prediction that metaphor interpretation involves the suppression of metaphor-inconsistent information, this mechanism seems to operate faster in the resolution of lexical ambiguity than in metaphor interpretation.

5 GENERAL DISCUSSION

The results of the present study offer support to the twofold hypothesis that metaphor interpretation involves enhancing relevant properties of the metaphor vehicle while suppressing irrelevant ones (Glucksberg & Keysar 1990; Carston 2002; Recanati 2004). Both superordinates and distinctive properties, which were, respectively, metaphor-inconsistent and metaphor-relevant properties in their contexts, were active up to 400 ms from the offset of the prime. However, at the longest delay, the level of activation of these associates was significantly different, with only the metaphor-relevant properties remaining active. In another lexical priming study using literal neutral contexts, we observed the opposite pattern of results, with superordinates remaining active up to 1000 ms but distinctive properties decaying between 400 and 1000 ms (Rubio et al. 2003). These different patterns of activation resulted in a crossover effect (Rubio 2005), which supports the view that distinctive properties were enhanced in metaphoric contexts where they were relevant for interpretation, whereas superordinates were actively suppressed given their inconsistency with the figurative interpretation of the nominal metaphor. Therefore, the present study offers support to previous studies of property activation in metaphor

interpretation (Gernsbacher et al. 2001; Glucksberg et al. 2001; McGlone & Manfredi 2001).

Regarding the issue of whether or not the processes involved in metaphor interpretation are inferential, it does seem that some of the nominal metaphors used in this study could be interpreted by purely local associative processes (Recanati 2004). For example, in interpreting 'Compared to the rest of the boys in the basketball team, John was a minnow', the activation of the information associated with BASKET-BALL TEAM would include information regarding the size of the players, which in turn would make more accessible the metaphor-relevant information of the lexical concept MINNOW. However, not all the metaphoric contexts could be interpreted in this way, some apparently requiring fully inferential processes in the construction of the ad hoc concept (Carston 2002; Wilson 2003). For example, in interpreting 'After six months without going to the barber, John was a lion', the hearer would draw the conclusion that John's hair was long on the basis of processing the first clause, which would increase the accessibility of the contextual assumption lions have manes in interpreting the nominal metaphor 'John was a lion'. Therefore, the implicature that John's hair was long, which is inferred from the first clause, would be used in deriving the figurative implication that John had a thick mass of hair, which would follow from John's belonging to the category LION* characterized by the property 'having a thick mass of hair'.

In order to support his claim that narrowing and loosening involve purely automatic, associative processes, Recanati could argue that in processing the latter example, activation would spread from BARBER to HAIR and from HAIR to MANE, making accessible the figurative interpretation of the nominal metaphor without any inferential process being involved. However, this would imply that any similar metaphoric context including the word 'barber' and using 'lion' as the metaphor vehicle would result in a similar interpretation, which is not necessarily the case. Consider the following example:

(4) After getting himself some new clothes and going to the barber, John was a lion.

Unlike in the previous example, the interpretation of 'John was a lion' in (4) would be based on the stereotypical view of lions as epitomizing pride and courage. Since there does not seem to be a possible chain of automatic associations between the schemata for buying clothes and going to the barber and John feeling self-confident, some inferential process must take place in interpreting the nominal metaphor along these lines (e.g. some backward inference about John being pleased with his new look). Therefore, the type of metaphoric contexts used in this experiment support the relevance-theoretic view of narrowing and loosening as inferential pragmatic processes rather than merely associative (Carston 2002; Wilson 2003).

On the other hand, the present results support Recanati's claim about the accessibility of the lexical concept in constructing the derived ad hoc concept (Recanati 1995, 2004; cf. Sperber & Wilson 2002). The superordinates used in this experiment, which were inconsistent with the figurative interpretation of the metaphor vehicle, were highly accessible up to 400 ms from its offset. However, I would not argue that the nominal metaphors were not interpreted until that point in processing, but rather that, because of their strong association with the prime, superordinates remained active past the point where they were discarded from interpretation. Given that the nominal metaphors in this experiment were novel, one may assume that they would have been perceived as figurative uses. The long activation of superordinates could therefore be understood along the lines put forward by Recanati (2001): although certain literal contextual assumptions would have been abandoned in interpreting the utterance figuratively, the corresponding conceptual properties would still be active, giving rise to a certain level of metaphor awareness.⁴

Although the activation pattern of superordinates can be explained following Recanati (2001, 2004), the results of the present experiment do not generally support his model of lexical processing in terms of accessibility. According to Recanati (1995, 2004), the selection of a conceptual representation is determined exclusively by its accessibility, so that the most highly accessible one will be selected for interpretation. In this study, the activation of metaphor-inconsistent and metaphor-relevant properties was not significantly different at 400 ms, with superordinates showing a slightly higher level of priming at that point in processing. Therefore, according to Recanati's model of lexical interpretation, metaphor-inconsistent properties would have been selected for interpretation at the intermediate delay, to be suppressed later on in processing. This interpretation of the results would also follow from the standard pragmatic view of metaphor,

⁴ The long activation of metaphor-inconsistent properties in the present study could be explained along the lines of 'the graded salient hypothesis' (Giora 2002 and references therein), according to which salient (coded, conventional) meanings are processed initially regardless of their contextual relevance. Although this hypothesis would predict the long activation of superordinates in metaphoric contexts as strong associates of the primes, it does not distinguish between a meaning component being highly accessible during processing and it being selected as part of the meaning communicated by use of a word (see Rubio 2005, forthcoming for discussion).

according to which the literal interpretation of a metaphoric expression is derived first, before looking for an alternative, figurative interpretation that is satisfactory in the context (Grice 1975, 1989; Searle 1979; for empirical evidence against this view see Glucksberg *et al.* 1982; Gildea & Glucksberg 1983; Keysar 1989).

For the present results to have fully supported Recanati's account of lexical interpretation (Recanati 1995, 2001, 2004), superordinates should have been active at the intermediate delay but at a lower level than distinctive properties, so that the latter would have been selected for interpretation while the former would have stood in conflict with interpretation. However, if the selection of meaning components was directed by considerations of relevance and not by accessibility alone (Carston 2002; Wilson & Sperber 2004), it would still be possible to argue that superordinates had been discarded from interpretation although still highly active at 400 ms. In this view, the selection of the most accessible properties would minimize processing costs, but the maximization of cognitive effects (especially the derivation of implications of the metaphoric meaning) would be the determining factor in selecting the relevant properties for interpretation. Thus, although superordinates would have been active at the intermediate delay because of their strong association with the prime concepts, only distinctive properties would have been selected for interpretation given that those properties would have given rise to greater cognitive effects when promoted to the status of content constitutive of the ad hoc concept which figured in the proposition expressed.

Regarding the operation of the mechanism of suppression in metaphor interpretation and disambiguation, the sustained activation of superordinates in metaphoric contexts seems to indicate that the suppression of these metaphor-inconsistent properties is different from the suppression of the irrelevant meanings of homonyms, which takes place 200-300 ms from the offset of the ambiguous prime (e.g. Tanenhaus et al. 1979, Seidenberg et al. 1982). It is therefore reasonable to conclude that the differential availability of the various meanings accessed in lexical ambiguity and novel metaphor interpretation has an effect on the operation of the mechanism of suppression in these two paradigms. Thus, although both disambiguation and novel metaphor interpretation involve dealing with two different concepts, in resolving a lexical ambiguity, various concepts are accessed from the start, although only one is selected in later stages of processing. In contrast, in processing a novel metaphorical expression, a single lexical concept is accessed initially, although an ad hoc concept is constructed on-line later on in processing. The inappropriate meaning of a homonym is

therefore suppressed in a context-sensitive, albeit effectively automatic, way given that the various meanings are available to the processor from the start. On the other hand, suppressing the literal meaning of a novel metaphor would require the operation of later, attentional processes since the alternative figurative meaning becomes available only during the process of understanding the utterance.

The idea that the inhibitory processes involved in metaphor interpretation are more demanding of attentional resources than those involved in disambiguation ties in with the idea that metaphorical language is 'special' (see Gibbs 1989 for discussion). A recurrent point in the literature on lexical processing is how, despite the initial activation of the various meanings of ambiguous words, people are not usually aware of having encountered an ambiguous word in a sentence (Tanenhaus *et al.* 1979; Gernsbacher 1990). On the other hand, we are usually aware of figurative language use, especially in the case of novel metaphors like the ones used in the present experiment. It is therefore possible that metaphor awareness is related to the fact that more attentional resources may be involved in metaphor interpretation as opposed to disambiguation, the processing of which is virtually automatic.⁵

In this view, metaphor interpretation would be special, but not necessarily more so than lexical pragmatic processes such as disambiguation or reference assignment, the difference lying merely in the degree of automatization of the cognitive processes involved. Lexicalized metaphors or metaphors that after frequent use have given rise to a second meaning of an expression (e.g. 'bulldozer' as in heavy machine and overbearing person) would be a case in point. If the association between a metaphor vehicle and one of its possible figurative interpretations is strengthened enough by frequent processing, the metaphor vehicle becomes polysemous. Therefore, it is possible that diachronically, the process of meaning construction involved in metaphor interpretation evolves through automatization into a process of meaning selection and so disambiguation (see Neely 1977 and Barsalou 1982 for a discussion of the automatization of controlled processes). It seems reasonable to assume that metaphor awareness would be lost at some point in that diachronic process of automatization of meanings, so the more familiar a metaphor, the less figurative (and probably evocative) it may seem.

⁵ Empirical evidence that reading literal and metaphorical expressions takes a comparable amount of time is usually taken to show that processing literal and figurative language takes the same amount of processing effort (see Gibbs 1984, 1989). However, these studies should be interpreted carefully as extra attentional resources can always reduce reading time while still adding to the processing effort account (see Rubio 2005 for discussion).

Regarding the role of suppression in lexical ambiguity and metaphor interpretation, it is possible that the literal interpretation of a metaphor vehicle does not need to be suppressed if the figurative one is made highly accessible by the preceding linguistic context, just as suppression does not need to operate on the inappropriate meaning of a homonym in a biasing context where the appropriate meaning has been lexically primed (Tabossi 1988; Simpson and Krueger 1991). However, this is not necessarily the case in all contexts. The results of the present study suggest that, in contexts where metaphor-relevant properties are not facilitated, suppression reduces the activation of metaphor-inconsistent properties of the vehicle concept.

I would therefore like to propose two intersecting continua that may help to delimit the extent to which extra attentional resources are required in processing a metaphorical expression. First, metaphor vehicles may range from novel to familiar to lexicalized, depending on how frequently they have been encountered previously (note that a lexicalized metaphor would be a polysemous word, with the two paradigms being compared touching at that point). This continuum would be related to the availability of the metaphorical meaning in the mental lexicon, with metaphorical meanings not being strongly associated to a novel metaphor but being accessed automatically if the vehicle is a lexicalized metaphor. Second, linguistic contexts may range from nonsensical to metaphoric to priming, depending on how much they facilitate the metaphorical interpretation of the expression. The first type of context would be comparable to ambiguous contexts (e.g. 'John is a banana' or 'John went to the bank' without any contextual cues that may help interpretation). Metaphoric contexts would be similar to disambiguating contexts in that the appropriate interpretation would be made accessible (e.g. 'In his flashy coat, John is a banana'), but without going to the extreme of priming it intralexically, in which case the context would be priming (e.g. 'In his long yellow coat, John is a banana').⁶

⁶ The career of metaphor model (Bowdle & Gentner 2005) posits that, as a metaphor vehicle becomes conventional, its processing changes from comparison to categorization (cf. Glucksberg & Haught 2006; Jones & Estes 2006). Unlike Bowdle and Gentner, I am not proposing that novel and familiar metaphors involve different modes of processing, but rather that constructing an ad hoc concept when interpreting a novel metaphor may be more demanding of attentional resources than retrieving a figurative meaning that has been lexicalized after frequent use. Another difference with the above studies is that I consider context to play a fundamental role in making a figurative interpretation more or less accessible and therefore I do not see metaphor aptness as a fixed value in the absence of a particular context.

The amount of attentional resources involved in interpreting a metaphorical expression would therefore be determined by the combination of these two factors: the degree of familiarity of the metaphorical interpretation and the strength of the contextual bias. It is obvious that the degree of familiarity of a literal expression and the bias of a literal context also combine to determine to some extent the amount of processing effort involved in understanding a literally used expression. However, I would still maintain that, at the lexical level, the availability of a lexically encoded concept as compared to an ad hoc concept constructed on-line makes metaphorical interpretation generally more dependent on context, and sometimes on attentional resources, than literal language interpretation.

APPENDICES: EXPERIMENTAL MATERIALS

Primes	Superordinates	Distinctive Features
Cactus	Plant	Spike
Lion	Animal	Mane
Slippers	Shoe	Comfortable
Skyscraper	Building	Tall
Lullaby	Song	Sleep
Dalmatian	Dog	Spot
Mercedes	Car	Expensive
Chair	Seat	Back
Champagne	Drink	Bubble
Breakfast	Meal	Morning
Cheetah	Cat	Fast
Sapling	Tree	Young
Woodpecker	Bird	Noise
Pacific	Ocean	Large
Rugby	Sport	Tough
Steel	Metal	Strong
Minnow	Fish	Small
Banana	Fruit	Yellow
Encyclopaedia	Book	Knowledge
Norway	Country	Cold

Appendix A: Primes and Targets

Appendix B: Metaphoric Contexts

- John doesn't like physical contact. Even his girlfriend finds it difficult to come close to him. John is a cactus.
- After six months without going to the barber, John was a lion.
- Mary is very materialistic. She is only interested in men who are rich. Her latest boyfriend is a Mercedes.
- John knew many people, but there were only a few friends he could lean on. A good friend is a comfy chair.
- John loved paddling his canoe through the steep canyon. He especially enjoyed rolling over in the white water of the rapids. The river was champagne.
- On the dunes someone had planted a few pine trees among the local plants. The pine trees were skyscrapers.
- Mary loved maths but this year the teacher was very boring. Every lesson was a lullaby.
- It was impossible to study at college during the maintenance work. The carpenter next door was a woodpecker.
- John spends four hours a day in the gym. His muscles are steel.
- Compared to the other boys in the basketball team, John was a minnow.

Nobody wanted to run against John at school. John was a cheetah. John likes to wear clothes that really stand out in the crowd. In his

- new coat, John is a banana.
- Even though she had never been to school, Mary was an encyclopaedia.

Mary had been sharing a flat with John for a long time. With him she felt at ease even in silence. John was a pair of old slippers.

- When he was a kid, John wasn't allowed to do many things on his own. His mother used to tell him that he was only a sapling.
- John was making a chocolate milkshake when the lid came off the blender. When his mother saw him, she said John was a Dalmatian.
- When Maria first came to England, she was very surprised that pubs closed at 11 pm. In Spain, closing time is breakfast.
- John and Mary have a new house with an amazing garden. Actually, one cannot see the end of it from the back door. Their garden is the Pacific.
- Things weren't going well for Mary. Her boyfriend had broken up with her the same week she had lost her job. Sometimes life can be a game of rugby.

Mary didn't like spending the night at her grandmother's. No matter how many blankets she would put on the bed, that attic room was Norway.

Acknowledgements

The study reported in this paper is part of the experiments completed by the author towards her Ph.D. degree at Cambridge University. This research was initially supported by an Arts and Humanities Research Board Postgraduate Award, and currently by a British Academy Postdoctoral Fellowship and a Marie Curie Outgoing International Fellowship (Project 022149—Inferential Processes in Language Interpretation). This paper was written at University College London and reviewed at Princeton University, where I am currently working on a postdoctoral project. I would like to thank Bart Geurts, Zachary Estes, Lara Jones and an anonymous reviewer for their comments and suggestions, which I think have improved the paper greatly. Thanks as well to Richard Breheny and John Williams for supervising the setting and analysis of my experiments. I would also like to thank Ira Noveck and Nausicaa Pouscoulous for their comments on an earlier version of this paper, as well as Robyn Carston and Sam Glucksberg for the valuable insight and support in all their comments.

PAULA RUBIO FERNÁNDEZ

Department of Phonetics and Linguistics University College London Gower Street London WC1E 6BT e-mail: paula@ling.ucl.ac.uk

REFERENCES

- Barsalou, L. W. (1982), 'Contextindependent and context-dependent information in concepts'. *Memory and Cognition* **10**:82–93.
- Barsalou, L. W. (1987), 'The instability of graded structure: Implications for the nature of concepts'. In U. Neisser (ed.), *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*. pp. 101–40. Cambridge University Press. Cambridge.
- Bowdle, B. & D. Gentner (2005), 'The career of metaphor'. *Psychological Review* **112**:193–216.

- Cacciari, C. & S. Glucksberg (1994), 'Understanding figurative language'. In M. Gernsbacher (ed.), *Handbook of Psycholinguistics*. Academic Press. pp. 447–77. San Diego, CA.
- Carston, R. (1996), 'Enrichment and loosening: Complementary processes in deriving the proposition expressed?' UCL Working Papers in Linguistics 8:61–88. Reprinted in (1997) Linguistische Berichte 8:103–27.
- Carston, R. (2002), Thoughts and Utterances: The Pragmatics of Explicit Communication. Blackwell. Oxford.

- Carston, R. (2007), 'How many pragmatic systems are there?' In M.-J. Frapolli (ed.), Saying, Meaning, Referring: Essays on the Philosophy of Francois Recanati. Palgrave. London, pp. 18–48.
- Gernsbacher, M. A. (1990), Language Comprehension as Structure Building. Lawrence Erlbaum. Hillsdale, NJ.
- Gernsbacher, M. A. & M. Faust (1991), 'The role of suppression in sentence comprehension'. In G. B. Simpson (ed.), Understanding Word and Sentence. Elsevier Science Publishers. North Holland.
- Gernsbacher, M. A., B. Keysar, & R. W. Robertson (1995), The role of suppression in metaphor interpretation. Paper presented at the annual meeting of the Psychonomic Society, Los Angeles, CA.
- Gernsbacher, M. A., B. Keysar, R. W. Robertson, & N. K. Werner (2001), 'The role of suppression and enhancement in understanding metaphors'. *Journal of Memory and Language* 45:433–50.
- Gibbs, R. W. (1984), 'Literal meaning and psychological theory'. *Cognitive Science* 8:275-304.
- Gibbs, R. W. (1989), 'Understanding and literal meaning'. *Cognitive Science* **13**:243–51.
- Gildea, P. & S. Glucksberg (1983), 'On understanding metaphor: The role of context'. *Journal of Verbal Learning and Verbal Behavior* **22**:577–90.
- Giora, R. (2002), 'Literal vs. figurative language: Different or equal?' *Journal* of *Pragmatics* **34**:487–506.
- Glucksberg, S. (2003), 'The psycholinguistics of metaphor'. *Trends in Cognitive Sciences* **7**:92–6.
- Glucksberg, S., P. Gildea, & H. B. Bookin, (1982), 'On understanding nonliteral speech: Can people ignore

metaphors?' Journal of Verbal Learning and Verbal Behavior 21:85–98.

- Glucksberg, S. & C. Haught (2005), 'On the relation between metaphor and simile: When comparison fails'. *Mind and Language* **21**:360–78.
- Glucksberg, S. & B. Keysar (1990), 'Understanding metaphorical comparisons: Beyond similarity'. *Psychological Review* **97**:3–18.
- Glucksberg, S., D. A. Manfredi, & M. S. McGlone (1997a), 'Metaphor comprehension: How metaphors create new categories'. In T. B. Ward, S. M. Smith, and J. Vaid (eds.), Creative Thought: An Investigation of Conceptual Structures and Processes. American Psychological Association. Washington, DC.
- Glucksberg, S., M. R. Newsome, & Y. Goldvarg (1997b), Filtering out irrelevant material during metaphor comprehension. Paper presented at the Meeting of the Psychonomic Society, Philadelphia, PA.
- Glucksberg, S., Newsome, M. R., & Y. Goldvarg (2001), 'Inhibition of the literal: Filtering metaphor-irrelevant information during metaphor comprehension'. *Metaphor and Symbol* **16**:277–98.
- Grice, H. P. (1975), 'Logic and conversation'. In P. Cole and J. Morgan (eds.), *Syntax and Semantics III: Speech Acts.* Academic Press. London.
- Grice, H. P. (1989), Studies in the Way of Words. Harvard University Press. Cambridge, MA.
- Johansson, S. & K. Hofland (1989), Frequency Analysis of English Vocabulary and Grammar based on the Lob Corpus: Vol. 1, Tag Frequencies and Word Frequencies. Clarendon Press. Oxford.
- Jones, L. L. & Z. Estes (2005), 'Roosters, robins and alarm clocks: Aptness and conventionality in metaphor

comprehension'. Journal of Memory and Language **55**:18–32.

- Keele, S. W. & W. T. Neill (1978), 'Mechanisms of attention'. In E. C. Carterette and M. P. Friedman (eds.), *Handbook of Perception IX*. Academic Press. New York.
- Keysar, B. (1989), 'On the functional equivalence of literal and metaphorical interpretations in discourse'. *Journal of Memory and Language* **28**:375–85.
- Marslen-Wilson, W. D. (1987), 'Functional parallelism in spoken wordrecognition'. Cognition 25:71–102.
- McGlone, M. S. & D. Manfredi (2001), 'Topic-vehicle interaction in metaphor comprehension'. *Memory and Cognition* **29**:1209–19.
- Neely, J. H. (1976), 'Semantic priming and retrieval from lexical memory: Evidence for facilitatory and inhibitory processes'. *Memory and Cognition* **4**:648–54.
- Neely, J. H. (1977), 'Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited capacity attention'. *Journal of Experimental Psychology* **106**:226–54.
- Neill, W. T., L. A. Valdes, & K. M. Terry (1995), 'Selective attention and the inhibitory control of attention'. In F. N. Dempster and C. J. Brainerd (eds.), *Interference and Inhibition in Cognition*. Academic Press. San Diego, CA.
- Neill, W. T. & R. L. Westberry (1987), 'Selective attention and the suppression of cognitive noise'. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **13**:327–34.
- Onifer, W. & D. A. Swinney (1981), 'Accessing lexical ambiguities during sentence comprehension: Effects of frequency of meaning and contextual bias'. *Memory and Cognition* **9**:225–6.
- Posner, M. I. & C. R. R. Snyder (1975), 'Attention and cognitive control'. In R. L. Solso (ed.), *Information Processing*

and Cognition: The Loyola Symposium. Lawrence Erlbaum. Hillsdale, NJ.

- Recanati, F. (1995), 'The alleged priority of literal interpretation'. *Cognitive Science* **19**:207–32.
- Recanati, F. (2001), 'Literal/nonliteral'. *Midwest Studies in Philosophy* **25**: 264–74.
- Recanati, F. (2004), *Literal Meaning*. Cambridge University Press. Cambridge.
- Rosch, E. & C. B. Mervis (1975), 'Family resemblances: Studies in the internal structure of categories'. *Cognitive Psychology* **4**:573–605.
- Rubio Fernández, P. (2005), 'Pragmatic processes and cognitive mechanisms in lexical interpretation: The on-line construction of concepts'. Ph.D. thesis, University of Cambridge.
- Rubio Fernández, P. 'Concept narrowing: The role of context-independent information'. *Journal of Semantics*, in press.
- Rubio Fernández, P., R. Breheny & Wei Lee (2003), 'Context-Μ. independent information in concepts: An investigation of the notion of "core features"". In F. Schmalhofer, R. M. Young, and G. Katz (eds.), Proceedings of EuroCogSci 03: The European Cognitive Science Conference 2003. Lawrence Erlbaum. London.
- Schvaneveldt, R. W. & D. E. Meyer (1973), 'Retrieval and comparison processes in semantic memory'. In S. Kornblum (ed.), Attention and Performance IV. Academic Press. New York.
- Searle, J. (1979), *Expression and Meaning*: Cambridge University Press. Cambridge, England.
- Seidenberg, M. S., M. K. Tanenhaus, J. L. Leiman, & M. Bienkowski (1982), 'Automatic access of the meanings of ambiguous words in context: Some limitations of

knowledge-based processing'. Cognitive Psychology 14:470-519.

- Shiffrin, R. M. & W. Schneider (1977), 'Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory'. *Psychological Review* 84:127–90.
- Simpson, G. B. & M. A. Krueger (1991), 'Selective access of homograph meanings in sentence context'. *Journal of Memory and Language* 30:627–43.
- Sperber, D. & D. Wilson (1986/1995), Relevance: Communication and Cognition. Blackwell. Oxford.
- Sperber, D. & D. Wilson (1997), 'The mapping between the mental and the public lexicon'. UCL Working Papers in Linguistics 9:107–25.
- Sperber, D. & D. Wilson (2002), 'Pragmatics, modularity and mindreading'. *Mind and Language* **17**:3–23.
- Swinney, D. A. (1979), 'Lexical access during sentence comprehension: (Re)consideration of context effects'. *Journal of Verbal Learning and Verbal Behavior* 18:645–59.
- Tabossi, P. (1988), 'Accessing lexical ambiguity in different types of sentential context'. *Journal of Memory and Language* 27:324–40.

- Tanenhaus, M. K., Leiman, J. M., & M. S. Seidenberg (1979), 'Evidence for multiple stages in the processing of ambiguous words in syntactic contexts'. *Journal of Verbal Learning* and Verbal Behavior 18: 427–40.
- Wilson, D. (2003), 'Relevance theory and lexical pragmatics'. *Italian Journal* of Linguistics/Rivista di Linguistica 15:273–91 (special issue on pragmatics and the lexicon).
- Wilson, D. & D. Sperber (2004), 'Relevance theory'. In G. Ward and L. Horn (eds.), *Handbook of Pragmatics*. Blackwell. Oxford. Longer earlier version published in (2002) UCL Working Papers in Linguistics, 14: 249–87.
- Winer, B. J. (1971), *Statistical Principles in Experimental Design*. McGraw-Hill. New York.
- Yee, P. L. (1991), 'Semantic inhibition of ignored words during a figure classification task'. *Quarterly Journal* of *Experimental Psychology* **43A**: 127–53.

First version received: 06.08.2006 Second version received: 29.01.2007 Accepted: 02.02.2007