

The Role of Experience in the Interpretation of Noun-Noun Combinations

Phil Maguire and Arthur Cater

Department of Computer Science, University College Dublin
Belfield, Dublin 4, Ireland.
phil.maguire@ucd.ie, arthur.cater@ucd.ie

Abstract. Previous studies [2], [6] have shown that combinations whose modifiers are typically associated with the instantiated relation are interpreted reliably faster than those whose modifiers are less frequently associated with the instantiated relation. Gagné and Shoben attributed this effect to the influence exerted by prior experience of the modifier. They proposed that speakers maintain relation type distributions regarding the frequency with which modifier nouns have been used with each of 16 possible relation types. However, others [4], [5], [7] have claimed that the differences in response times observed in Gagné and Shoben’s study may have arisen due to the nature of the concepts being combined. We contrasted these views by investigating whether these differences persisted when the constituent concepts of a combination were presented individually. We presented noun pairs without a modifier-head syntax; as a result interpretation could not be influenced by previous experience of how those nouns had been used as modifiers in the past. The results revealed that differences in response time remained using either method of presentation. This finding is problematic for the CARIN theory and, as a result, we consider other factors that might influence the difficulty of linking two arbitrary concepts.

1 Introduction

In everyday conversation, noun-noun compounds (also termed conceptual combinations) are frequently used in order to express new ideas and to encapsulate novel objects and experiences (e.g. beer headache, peasant dance). Compounding is a useful practice in that it greatly enhances the flexibility of language as well as increasing efficiency in communication. Although people have a well developed means of understanding these novel compounds, the comprehension process is often not trivial, requiring an understanding of the speaker’s communicative goals and a representation of the situation at hand as well as detailed world knowledge. Accordingly, the study of conceptual combination is important because of the way in which it is intimately associated with the generativity and comprehension of natural language. In English, where compounding is particularly productive, combinations consist of a modifier followed by a head noun. Usually, the head noun denotes the main category while the modifier implies a relevant subcategory or a modification of this set’s typical members. In this way, a *mountain flower* is interpreted as a type of flower, and more particularly as one which is located in mountains. In order to understand a combination like this, people

first have to deduce the most likely way in which *mountain* could be intended as a modification of the concept *flower*. Once a link between the two concepts has been definitively established, the combined concept can then be elaborated upon and further inferences can be made (e.g. *mountain flowers* are likely to be wild).

Gagné and Shoben demonstrated in [2] that the more frequently a relation is associated with the modifier noun of a combination, the easier it is to judge whether a combination involving that relation is sensible or not. The effect itself seems relatively intuitive. Consider the modifier *chocolate*: combinations in the form *chocolate X* can be most easily interpreted as *X <made of> chocolate* seeing as the most common instantiation of *chocolate* as a modifier involves the compositional <H *made of* M> relationship. In this way, even the combination *chocolate train* is more likely to be interpreted as a train made of chocolate than as a train containing chocolate, despite the latter perhaps being the more plausible. It also seems reasonable that combinations in which modifiers are used in an atypical fashion might prove more difficult to interpret: *chocolate magazine* prompts a momentary hesitation during which the most obvious instantiation of *chocolate* must be suppressed. Based on these principles, Gagné and Shoben attributed differences in response times observed during a sensibility judgement task to prior experience of the modifier [2]. Their Competition Among Relations in Nominals (CARIN) theory proposes that speakers maintain distributions recording how frequently (in relative terms) nouns have combined as modifiers using each of 16 possible relation types outlined therein.

Despite the apparent reasonableness of this theory, others have questioned whether prior experience is the central contributing factor towards the observed phenomenon [4], [5], [7]. Murphy argues in [5] that the storage of arbitrary relations for each noun, as proposed by CARIN, seems unconvincing mainly because the 16 relation types provided are ambiguous and nondescriptive. In effect, relations are extremely varied and detailed and as a result, classifying a relation under the CARIN taxonomy abstracts it to such a level that such information becomes useless. The focus of our study is thus on contrasting the CARIN theory with the alternative view that differences in ease of interpretation arise naturally and are not caused by prior experience of the modifier. The following experiment accomplishes this by examining the process of relation selection between noun concepts in the absence of a modifier-head syntax.

1.1 Overview of Rationale

The most unambiguous method by which to address the role of prior experience in conceptual combination would be to investigate cases in which no such experience is available. If it were practical, one might recruit speakers of a combination-free dialect as participants: speakers encountering combinations for the first time would be required to interpret them in the absence of any statistical knowledge, thereby providing an ideal condition for investigating CARIN's premises. Unfortunately, given the ubiquity of conceptual combination, virtually every language contains combinations of some form, thereby rendering this paradigm infeasible. Another approach would be to present the concepts as images, thus avoiding the activation of linguistic knowledge. Unfortunately, the use of images is unreliable and picture recognition is greatly influenced by canonicity and view specificity. Furthermore, some concepts cannot be represented pictorially in their prototypical form: a representation of fruit would in-

variably refer to one particular type of fruit. Similarly, conveying abstract concepts like *dilemma* or *justice* would be even more problematic.

In order to sidestep these problems while simultaneously suppressing the activation of statistical knowledge, an alternative method of presentation was adopted. This paradigm was based on the assumption that a noun concept cannot be interpreted as a modifier in the absence of a head on which it can act. Furthermore, statistical knowledge regarding relation type preference is irrelevant when the noun concept is not acting as a modifier. As a result, CARIN-style statistical knowledge can have no influence on the comprehension of an isolated noun concept. For example, in reading the noun *chocolate* on its own, only the material noun interpretation is activated. Its alternative interpretations such as *<made of chocolate>* or *<containing chocolate>* appear only in combination with an appropriate head. As a result, the distribution of *chocolate*'s relation type frequencies cannot influence the comprehension process in this case.

We therefore decided to present noun concept pairs as individual nouns. This ensured that the nouns would not be interpreted as constituents of a combination, and hence that statistical knowledge regarding the modifier's relation type preference would not be activated. By controlling the influence of past experience in this way, we were able to investigate the nature of relation selection when that process is based solely on the conceptual content of the constituent nouns.

2 Experiment

The experiment was designed to eliminate any influence of modifier history. Presenting nouns outside a combinational syntax obliged participants to relate the noun pairs by virtue of their semantic representation only. According to the CARIN theory, differences in response time should only be apparent in cases where statistical knowledge about modifier relation type frequencies is activated. As a result we expected that the differences in ease of interpretation found by Gagné and Shoben in [2] would disappear given this novel method of presentation. By investigating whether differences in ease of interpretation are mitigated when noun concept pairs are interpreted independently, we hoped to corroborate the claim that such differences arise due to statistical knowledge.

Two separate conditions were identified, both involving the same experimental materials. In the combined concept (CC) condition, participants were presented with noun pairs in the typical combinational format and were required to judge the sensicality of these combinations. In the independent concept (IC) condition, a different set of participants was presented with the same nouns in isolation and these participants were asked to decide whether a sensible linking relation could be found. According to CARIN, the interpretation of a novel noun-noun combination occurs when speakers identify a relation between the modifier and the head noun. Thus, as participants in the IC condition were required to search for a sensible relation linking the nouns, the participants in both conditions were effectively carrying out the same ultimate task, albeit in different ways. In the CC condition, participants could exploit statistical knowledge as CARIN supposes; in the IC condition, they could not.

2.1 Method

Participants Forty first-year undergraduate students from University College Dublin participated in the study for partial course credit. All were native English speakers.

Materials For the sake of comparison, we selected a subset of the materials used by Gagné and Shoben in [2] for our experiment. In their study, materials were divided into three categories, namely HH, HL and LH conditions. For these condition labels, the letters H and L refer to the frequency of the instantiated relation (High or Low), with the first letter denoting how frequently that relation is associated with the modifier and the second denoting how frequently it is associated with the head. Thus *mountain magazine* (a magazine <about> mountains) is considered to be an LH combination since the <about> relation is low-frequency for the modifier *mountain* but high-frequency for the head *magazine*.

We selected a sample of 10 materials from each of the HH, HL and LH categories for the purposes of the experiment. These materials were selected based on several criteria. Any of Gagné and Shoben’s modifiers that were adjectives (e.g. *thermal, historical*) were excluded from consideration. We also ignored combinations in which the modifying noun was actually intended in a plural sense despite being denoted as singular (e.g. *servant language* is a language used by servants in general and not one servant in particular). In a re-analysis of the original CARIN study, Wisniewski and Murphy [7] suggested that the plausibility and familiarity of Gagné and Shoben’s materials may not have been properly controlled, thus confounding response times for the various conditions. Indeed, many of Gagné and Shoben’s materials are quite bizarre (e.g. *olive area, cooking hole*). In order to account for this inconsistency, two independent judges were asked to decide which of the eligible materials from the appendices of [2] were truly sensical and which ones were not. Any of the combinations judged as non-sensical by either of the judges was excluded. Following this, 10 materials were randomly selected from the remaining items in each category.

Several measures were taken to ensure that noun pairs presented in the IC condition would not be interpreted as a combination. Firstly, the presentation of each noun in a pair was separated by a 1,000ms visual and phonological mask consisting of a blank screen containing the “+” symbol. This was intended to prevent participants’ reading of the nouns coalescing into a combinational syntax within their phonological loop. As a further measure, we took advantage of the fact that modifiers in the English language are nearly always singular: having a plural in the modifier slot violates the (weak) constraints for combinational syntax. Thus when two nouns are presented together and the first is plural, the phrase can typically only be interpreted as two separate nouns and not as a combination (e.g., the two words *mountains + cloud* cannot be read as a single unit). We therefore pluralized a large portion of the filler modifiers in the IC condition (e.g. *dogs + vet, sandwiches + filling*). The presence of these pluralized fillers was intended to prevent participants from adopting the habit of viewing the two nouns as a combination, which might have supported their search for a linking relation.

The experimental materials used in both the CC and the IC conditions were identical. The method of presentation was different only in that the mask screen for the IC group contained the “+” symbol, whereas for the CC group it was completely blank. Aside from the instructions, the only other difference between the two conditions was the plurality of the filler modifiers. In the IC condition, we used 30 sensical combina-

tions with pluralized modifying nouns, 10 nonsensical combinations with singular modifying nouns and 10 nonsensical combinations with pluralized modifying nouns. In the CC condition, all of the corresponding fillers had singular modifiers.

Design A mixed 2 X 3 factorial design was used, with the two conditions of presentation as the between-participants measure (IC and CC), and Gagné and Shoben's [2] three conditions of relation type frequency as the within-participants measure (HH, HL and LH). The dependent measures were response time and accuracy rate, which were used to infer ease of interpretation. Participants were randomly assigned to the between-participants conditions, with a total of twenty in each. Each participant was presented with 80 stimuli, comprising the same set of 30 experimental stimuli for both between-subject conditions and the 50 filler items.

Procedure Participants sat in front of a computer screen and placed the index finger of their left hand on the F key of the computer keyboard and the index finger of their right hand on the J key. The participants in the CC condition were instructed that they would be presented with a series of concept combinations for which they had to make sensicality judgements and that the constituent nouns of the combinations would be displayed one after the other. In contrast, participants in the IC condition were instructed that they would be shown two concepts, and that they had to decide whether one concept could be combined with the other in a sensible manner. Both sets of participants were instructed to press J for sense and F for nonsense and emphasis was placed on the fact that they should only press F if the item was truly incomprehensible. In both conditions, each word was displayed by itself for one second in the centre of the screen, separated by the mask screen lasting for one second. After viewing the second word in the trial, participants had to make a sensicality judgement by pressing the appropriate key. The same materials were presented in both conditions and so the modifying noun was always presented first, although this was not made explicit.

Participants were initially given a short practice session where feedback was given regarding their judgements. The aim of this practice was to set a reliable threshold for sensicality and also to familiarize them with the nature of the task. Participants in the IC condition were shown pairs of concepts with pluralized modifying nouns (e.g. *tomatoes* + *sandwich*, *dogs* + *allergy*). After making a sensicality judgement participants were then informed whether the two nouns could be related and if so, the nature of the relation (e.g. "an allergy caused by dogs"). Similarly, in the CC condition, participants were shown concept pairs in a combinational format and after making a sensicality judgement, were shown how the combination was sensical or otherwise. Upon completing this practice session, participants were informed that they were beginning the main part of the experiment. The materials were then presented in a random order to each participant.

2.2 Results and Discussion

A total of 20.5% of trials were omitted from the analysis of the results, 15% in the CC condition and 26% in the IC condition. In 11.9% of CC trials and 23% of IC trials the incorrect response was given and hence these data were excluded. Responses were also eliminated if they were deemed unreasonably fast (0.1% of trials < 400 ms in the CC

condition) or slow (1.7% > 4000 ms in the CC condition, 2.5% > 8000ms in the IC condition). After this initial elimination process, any remaining response times which were more than three standard deviations outside each participant's mean were also excluded. This removed another 1.3% of trials in the CC condition and 0.5% in the IC condition. In the CC condition, the mean response time was 1,093; 1,145 and 1,254 ms for the HH, HL and LH conditions respectively while in the IC condition, the mean response time was 1,853; 1,873 and 2,213 ms. Response time was analysed using a 2 X 3 repeated measures ANOVA, with participants as a random variable, presentation method as a between-participants variable and relation type frequency as a within-participants variable. The ANOVA revealed a main effect of relation frequency, $F(2, 76) = 7.89, p < .01$. There was also a significant main effect of presentation method: response times were significantly longer in the IC condition than in the CC condition, $F(1, 38) = 15.81, p < .01$. Contrary to our hypothesis, there was no significant interaction between method of presentation and relation type frequency, indicating that the influence of relation frequency was not affected by method of presentation, $F(2, 76) = 1.89, p = .16$. The mean accuracy rates for the HH, HL and LH in CC condition were .93, .87 and .76 respectively while in the IC condition the mean accuracy rates were .83, .76 and .64. A second repeated-measures ANOVA revealed a main effect of relation frequency, $F(2, 76) = 26.21, p < .01$, and of presentation method, $F(1, 38) = 21.17, p < .01$. Once again, there was no significant interaction between accuracy and method of presentation, $F(2, 76) = .20, p = .82$, indicating that the pattern of accuracy rates was similar for both methods of presentation.

These results show significant differences in response time and accuracy rates between the high and low frequency modifier conditions for both the CC and IC conditions. The absence of an interaction between method of presentation and modifier relation frequency contradicts our hypothesis that eliminating the influence of statistical knowledge would mitigate differences in response time. It also suggests that the factors influencing response time and accuracy rates are not related to past experience of the modifier. Given that relation selection in the IC condition could only be carried out based on the semantic representation of the constituent concepts, it appears that the variations in response time, equally evident in both conditions, can only be accounted for by naturally arising differences in the ease of combining an arbitrary pair of noun concepts.

2.3 Correlation analyses

In order to determine the contribution of modifier influence towards the overall variance in response time, we obtained a correlation between response time and relation type frequency. Each of the materials was assigned a value corresponding to the relative frequency with which the modifier was associated with the instantiated relation. We used the same frequencies as Gagné and Shoben [2], which they derived by pairing 91 heads with 91 modifiers. After rejecting uninterpretable combinations, they analysed the relative frequencies with which the modifiers combined using each of the possible relation types. For example, *mountain* is typically interpreted using the <located> relation (e.g. *mountain cloud*) and according to Gagné and Shoben's frequencies, the relative frequency of this relation for *mountain* is .82. Using these values, we obtained correlations between response time and relation frequency and between accu-

racy rate and relation frequency. Neither of the correlations was significant for either the CC ($r = -.11, p = .58, r = .31, p = .09$) or the IC conditions ($r = -.24, p = .21, r = .28, p = .14$), challenging the notion that the relation type frequency of the modifier is an important factor in the interpretation process.

We also examined how well plausibility and familiarity predicted response time relative to modifier relation frequency. For these correlations we used the ratings provided by a group of 30 participants in Wisniewski and Murphy's re-analysis [7] of Gagné and Shoben's study [2]. The correlations between mean response time and familiarity and plausibility were highly significant for the CC ($r = -.61, r = -.56$ respectively) and the IC conditions ($r = -.57, r = -.58$). The correlations between accuracy rate and familiarity and plausibility were also highly significant for both the CC ($r = .69, r = .66$) and the IC conditions ($r = .72, r = .73$).

The fact that there was no significant correlation between modifier relation frequency and the dependent variables is a strong indicator that past experience of the modifier does not have a large influence on ease of interpretation. On the other hand, the high correlations involving plausibility and familiarity suggest that these factors account for a far greater portion of the variance in response time: stimuli referring to a more plausible and familiar concept were interpreted reliably faster and more accurately regardless of being presented as a combination or otherwise. These findings are consistent with those of Wisniewski and Murphy's in [7], which revealed that familiarity and plausibility are the strongest predictors of response time. Because these variables relate to the combined concept itself, they are thus independent of the method of presentation and this might explain how comparable differences were observed in both the CC and IC conditions.

In a further analysis, we investigated how well the ease of finding a relation between two separate noun concepts could predict the ease of interpreting the same concepts presented as a combination. In order to do this we correlated the two sets of dependent variable values from the IC and CC conditions, as both contained the same stimuli. The correlation between accuracy rates in the IC and CC conditions was highly significant, $r = .78$. The correlation between response times in both conditions was also significant, $r = .44$. Ignoring two pairs of stimuli for which the correct response was elicited less than 25% of the time in the IC condition, the correlation between both sets of response times increased to $.52$, which is substantial given the inherent variability of response times in general. These correlations indicate that noun pairs that were easy to interpret as a combination were also easy to relate when presented as individual noun concepts. Similarly, combinations that were often misjudged as nonsense when presented as a combination were also frequently misjudged when presented as two separate concepts. The high correlation between response times in both conditions suggests that the search for a linking relation accounts for a significant portion of the variance in the ease of interpreting a combination and that relation identification forms a fundamental part of the comprehension process.

At first blush, the strong correlation between familiarity and the dependent variables would seem to suggest that the frequency with which speakers are exposed to a certain combination directly influences how difficult that combination is to interpret. Certainly, combinations that are encountered very frequently can be stored as single entries in the lexicon, thereby obviating the combination process. However, a pair of one-tailed z-tests revealed no significant differences between the familiarity correlation

coefficients in the CC condition and those in the IC condition ($z = -.23$, $p = .41$, $z = -.22$, $p = .41$ for response time and accuracy correlations respectively). As materials in the IC condition were not presented as a combination, this suggests that familiarity of the combinational phrases per se was not responsible for the high correlations. As a result, it seems unlikely that participants were recalling previous encounters with the stimuli, which indeed were relatively novel.

Given this finding, the influence of familiarity must be due to factors other than memory retrieval. One possible explanation is that the familiarity ratings reflect the familiarity of the *referent concept* and not of the phrase itself: such familiarity would influence the ease of interpretation of a compound noun phrase whether it was presented as a combination or not. Another possible explanation is that familiarity covaries with plausibility: combinations which refer to more plausible concepts will happen to be encountered more frequently and will thus be rated as more familiar. This possibility is supported by a surprisingly high correlation of .94 between the plausibility and familiarity ratings for our materials. As a result, even though familiarity might appear to influence response time, this variable might only be a reflection of the referent's plausibility, this being the fundamental factor affecting the ease of interpretation.

2.4 Method of Interpretation

Given our results, one might propose that despite our efforts, participants in the IC condition were somehow interpreting the modifying nouns in a modifier sense. Several findings cast doubt on this possibility. Firstly, a main effect of presentation condition was observed and response times in the IC condition were significantly longer than those in the CC condition. Moreover, accuracy rates were significantly lower. If participants had been processing the word pairs as combinations then overall differences in response times or accuracy rates would not have been expected. One might argue that statistical knowledge about how a noun concept can be used as a modifier is still activated even when that noun is not being interpreted in a modifying role. In order to investigate this possibility we ran an analysis of the differences between response times to stimuli in the high modifier frequency conditions (HH and HL) and the sensical pluralized fillers. According to CARIN, combinations in the HH and HL conditions should benefit from statistical knowledge whereas pluralized modifiers should have no history because plural nouns like *tomatoes* are almost never used as modifiers. Despite this, we discovered that participants interpreted the pluralized modifiers significantly faster than the high modifier frequency stimuli, $t(19) = 2.89$, $p < .01$. This is the opposite to what would have been expected had participants been benefiting from the availability of statistical knowledge. It therefore suggests that the IC materials were evaluated based solely on the semantic representation of the individual noun concepts.

Despite the fact that the method of presentation did not affect the relative differences in response times between the relation type frequency conditions, it is worth noting that materials in the IC condition were interpreted reliably slower than those in the CC condition (an average of about 1800ms as opposed to 1200ms). This disparity indicates that the presence of a combinational syntax greatly enhances the fluency of interpretation. Although the effect may have been partially due to the unnaturalness of the task, the longer response times in the IC condition suggest that the syntactical constraints imposed by having a designated modifier and a designated head are important for fa-

cilitating relation selection: knowing which noun is acting as a modifier greatly speeds up the interpretation process. Without the clue afforded by syntax, participants in the IC condition may have felt the need to select among a considerably greater number of possible relationships, in some of which the first of the two concepts filled a head slot and the second filled the modifier slot. Furthermore, the presence of a designated modifier may have streamlined the interpretation process as the head concept is typically evaluated in light of the modifier, thereby obviating the full activation of both concepts. In this way, the most relevant features regarding modification can be quickly identified while redundant information can be avoided, a process which may not have been possible in the IC condition. Considering as an example the stimulus *gas lamp*, participants in the CC condition evaluated the concept *lamp* in the context of it being related to *gas*, thus arriving directly at the referent concept. On the other hand, participants in the IC condition may have activated both concepts more fully before searching for an appropriate linking relation, thereby triggering a representation of the prototypical electrical lamp and then being forced to reconsider.

2.5 Implications for the CARIN theory

Various probabilistic models of human language comprehension have been proposed in the past, based on the idea that probabilistic information about words, phrases and other linguistic structure is represented in the minds of language users and plays a role in language comprehension. Indeed, experiments related to general statistical language models show that humans are in fact very good predictors of word usage (see [3]). This would seem to suggest that prior experience as well as something akin to frequency distributions could indeed be a factor in human language processing. Although the current findings challenge the CARIN theory, they do not deny such a possibility. In certain circumstances, speakers may well be aware of the more typical usages of a modifier and this might affect how a combination is likely to be interpreted and hence the level of context required to ensure the correct interpretation. However, as an overall theory of conceptual combination, CARIN suffers from several limitations. The relation types suggested by the theory are somewhat arbitrary and ambiguous; many relations cannot be satisfactorily classified under CARIN's taxonomy while others can be placed into several categories. In addition, these relation categories sacrifice much of the detail and natural variation present, therefore rendering such labels uninformative (see [1]). However, CARIN's greatest inadequacy is that it fails to offer any explanation as to how the correct relation is eventually selected. Relation type frequencies alone can never suffice, as the interpretation process unavoidably requires the detailed consideration of both constituent concepts. For these reasons, it seems reasonable to accept that differences in ease of interpretation are predominantly dependent on the properties of the concepts being combined and indeed, this is the most reasonable explanation for our results.

3 Conclusion

Experiment has shown that the differences in response time observed by Gagné and Shoben in [2] are not eliminated by presenting stimuli without a modifier-head syntax. Because the experiment was designed so that participants would seek relations between concepts rather than words, any knowledge about the properties of words was rendered useless. As a result, our findings fail to support a central tenet of the CARIN theory, namely that ease of interpretation is significantly influenced by the combinational history of the modifier word. Correlations between modifier relation frequency and the dependent variables were not reliable. Conversely, correlations between plausibility and familiarity and the dependent variables proved highly significant, suggesting that these are far better predictors of ease of interpretation. The fact that the presentation of a combination as two individual noun concepts had no reliable influence on the relative differences in response times suggests that any factors that do influence interpretation are not linked to the combinational format itself but rather to the properties of the constituent nouns being combined. As a result, greater emphasis should be placed on understanding how the properties of multiple concepts are reconciled rather than focusing on how those words have been used in combinations in the past.

Importantly, this study has demonstrated that concept pairs can be related without a modifier-head syntax. The results revealed that the time taken to relate these concepts in isolation was a strong predictor of the time taken to interpret the corresponding combination. This highlights the important role of relation selection in the interpretation process and indicates that much of the variance in ease of interpretation can be accounted for by the complexity of identifying an appropriate relation. For this reason, future research should seek to explain both the process by which a linking relation is determined and the way in which a combinational syntax facilitates this process.

References

1. Devereux, B. & Costello, F.: Assessing the taxonomic approach to noun-noun compounds. Proceedings of the European Cognitive Science Conference (2003)
2. Gagné, C. L. & Shoben, E. J.: Influence of thematic relations on the comprehension of modifier-noun combinations. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 23 (1997) 71-87
3. Jurafsky, D.: Probabilistic Modeling in Psycholinguistics: Linguistic Comprehension and Production. In Bod, R., Hay, J. & Jannedy, S. (eds.) *Probabilistic Linguistics*. Cambridge, MA: MIT Press (2003)
4. Maguire, P. & Cater, A.: Interpreting noun-noun compounds with rare modifiers. Proceedings of the Seventh Conference of the German Cognitive Science Society (2005)
5. Murphy, G. L.: *The Big Book of Concepts*. Cambridge, MA: MIT Press (2002) 463-464
6. Storms, G. & Wisniewski, E. J.: Does the order of head noun and modifier explain response times in conceptual combination? *Memory and Cognition* (in press)
7. Wisniewski, E. J., & Murphy, G. L.: Frequency of relation type as a determinant of conceptual combination: a reanalysis. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 31 (2005) 169-74