

Turnip Soup: Head Noun Influence in the Comprehension of Novel Compounds

Phil Maguire (phil.maguire@ucd.ie)

Arthur Cater (arthur.cater@ucd.ie)

Department of Computer Science, University College Dublin
Belfield, Dublin 4, Ireland

Abstract

Although previous studies of noun-noun compounding have shown that past experience of a modifier influences the comprehension of combinations involving that modifier (Gagné & Shoben, 1997; Storms & Wisniewski, in press), no similar head noun influence has been found. In this paper we describe an experiment that investigates the influence exerted independently by head nouns that are particularly suggestive of the instantiated relation. Our results indicate that when suitable head nouns are examined, they exert a similar influence on comprehension difficulty to that exerted by modifiers. Combinations with biased heads were interpreted reliably faster than those with neutral heads and the head's strength ratio contributed significantly in predicting sensibility judgment times in a regression analysis. However, while a facilitating effect was observed for compounds using a relation highly typical of the head, no inhibiting effect was evident for relations that contradicted the head's bias. We discuss the implications of these findings and suggest some refinements to Gagné and Shoben's CARIN model.

Introduction

The combination of two words is a technique commonly adopted by speakers in order to refer to novel concepts (e.g. *holiday tension*, *picnic bee*). Although people have a well developed means of understanding these novel compounds, the associated comprehension process is not trivial, requiring many layers of understanding. Accordingly, the study of conceptual combination is important, both because it is intimately associated with the generativity and comprehension of natural language and because it is important for understanding how people represent concepts.

In English, a language in which 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 the set's typical members. In this way, a *cereal box* is interpreted as a particular type of box, and more precisely as one which contains cereal as opposed to other types of boxes. However, in order to understand a combination like this, people first have to be able to relate the two concepts in a meaningful way. Gagné and Shoben's (1997) Competition Among Relations In Nominals (CARIN) theory focuses predominantly on the relation between the constituent nouns in a combination. This theory maintains that there is a fixed, relatively small taxonomy of standard relations that can be used to link the

modifier and head noun concepts and that people possess statistical knowledge about how often each relation has been used with a modifier concept in the past. As a result, the most available standard relation is the one most frequently used to interpret other compounds containing that same modifier. For instance, the modifier *mountain* is most often associated with the <head LOCATED modifier> relation thus making the combination *mountain stream* easier to interpret than *mountain magazine* which uses the <head ABOUT modifier> relation.

The emphasis placed on the primacy of the modifier is an important feature of CARIN that distinguishes it from previous theories of conceptual combination. Gagné and Shoben (1997) provided empirical support in favor of this primacy. They found that the modifier's relation type frequency distribution influences the ease with which a combined concept can be interpreted while that of the head does not. Subsequent studies have replicated this effect in Indonesian (Storms & Wisniewski, in press) and in French (Maguire & Cater, 2004), languages in which the order of the head and the modifier is reversed. These studies showed that combinations involving a relation used frequently with the modifier were easier to interpret than combinations involving a less frequent relation, while no effect was observed regarding the relation type distribution of the head.

Although Gagné and Shoben (1997) pointed out that CARIN does not exclude a role for the head noun, the theory fails to elaborate on what form such a role might take. Despite the fact that no head influence was observed in their study, we speculated that a more sensitive experiment using carefully constructed stimuli would clarify the issue of the head's importance. Accordingly, the focus of our study was whether, under certain conditions, the relation type distribution of the head influences the interpretation process in the same manner as the modifier. We hypothesized that such an influence would be most apparent when the head is strongly biased towards the instantiated relation, as the requisite relation of such combinations is particularly available.

Overview of Rationale

The following experiment was designed in order to determine whether the relation type distribution of the head noun influences conceptual combination when the influence of the modifier is held constant. We examined whether combinations with heads that are particularly biased towards

the instantiated relation are easier to understand than combinations without this bias. In addition, we investigated whether combinations involving relations that contradict the head’s bias are more difficult to understand.

In order to carry out the study we needed to find head nouns that were reliably biased towards one relation type (e.g. *soup* is generally modified by a food substance such as *mushroom*). Consequently, we needed to find a way in which we could objectively measure the frequencies with which constituent nouns are associated with each relation type. Previous studies involving relation type distributions have faced the same problem and have tackled the issue using several different methods. Gagné and Shoben (1997) used a sampling technique in which they took 91 heads and 91 modifiers, paired them together and derived distribution frequencies for each based on the 3,239 sensible combinations that emerged. However, it is doubtful that a sample of this size is broad enough to accurately reflect the relation type distribution associated with a given noun. Additionally, this method fails to account for the fact that common combinations are encountered regularly, while exposure to rarer combinations may be very infrequent. It also ignores the way that different heads often combine with disparate sets of modifiers and vice versa. For this reason, methodologies relying on a constrained set of nouns will cause distortions in the overall relation type distributions. Furthermore, the unnaturalness of arbitrary pairings has the potential to exaggerate the incidence of rarer relations.

Storms and Wisniewski (in press) suggest an alternative method for deriving relation distributions that avoids the problems of unnatural context and type-token frequencies. In their study participants were provided with heads and modifiers independently and were asked to generate up to ten possible combinations containing each noun. After classifying all the generated combinations in terms of Gagné and Shoben’s relation categories, the relation frequencies were derived through counting. However, this technique is also prone to inaccuracies because participants are predisposed to select the most obvious combinations first, leading to much repetition. Hence, the resultant set of combinations fails to reflect the full range of combinatorial possibilities, with the result that the rarer relations are understated.

In order to avoid these problems and derive accurate relation type distributions, we used the British National Corpus (BNC) World Edition, and extracted all the noun-noun compounds from the 100 million words contained therein. This was carried out using the Gsearch tool (Corley et al., 2001), a chart parser that detects syntactic patterns in a tagged corpus via a user-specified context-free grammar. In order to determine relation type frequencies for heads that we suspected might show a skewed distribution, a random sample of 100 combinations containing the relevant head was extracted. These were then classified in terms of Gagné and Shoben’s (1997) relation categories.

Experiment

While previous studies have not found any influence exerted by the relation type distribution of the head noun, this may have been because the materials used were not sufficiently biased towards the dominant relation. Accordingly, we consciously sought out heads with the most skewed distributions and were able to verify these biases using the BNC frequencies.

The experiment was designed so that response times for combinations involving highly biased or “predictive” heads (P_{BIAS}) could be compared with response times for combinations using the same modifier and relation but with a less predictive head ($P_{CONTROL}$). In a second group of conditions we also compared response times between combinations where the relation contradicted a highly predictable relation suggested by the head (N_{BIAS}) and combinations using the same modifier and relation with an unbiased head ($N_{CONTROL}$). This second group of conditions was included so that we could examine whether comprehension becomes more difficult when an extremely atypical and unexpected relation is instantiated. All four conditions are illustrated in Figure 1 with relevant examples.

The focal difference between the P_{BIAS} and $P_{CONTROL}$ conditions was thus the relation type distribution of the head, and likewise for the N_{BIAS} and $N_{CONTROL}$ conditions. The modifier and relation were maintained across the two sets of conditions so that the influence of the relation type distribution of the modifier would be constant, and hence the influence exerted by the head could be examined independently. When selecting the materials we were also careful to avoid using any modifiers that were themselves particularly biased in case an extremely predictive modifier would overwhelm any influence exerted by the head (e.g. time periods, geographical locations, substances).

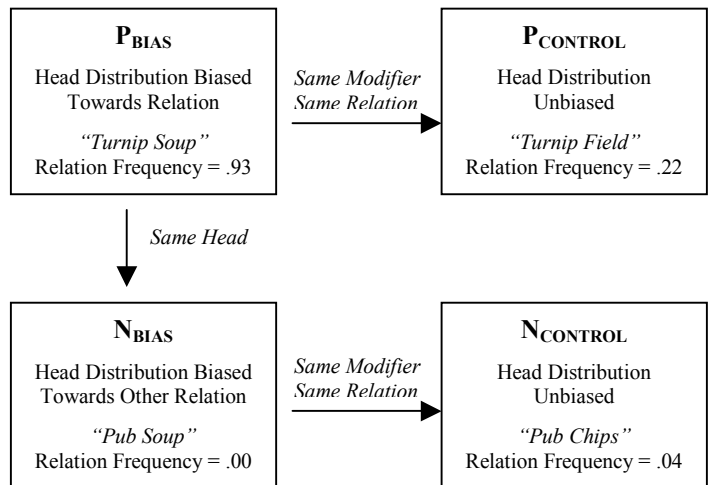


Figure 1: Experimental conditions.

Method

Participants Thirty-six first year undergraduate students from University College Dublin participated in the study for partial course credit.

Materials Twenty-five combined concepts were generated for each of the four conditions as well as 100 nonsensical filler items. Each of the sensible combinations was ascribed to one of the 16 relation types in the CARIN taxonomy. The 25 predictive heads used in the P_{BIAS} and N_{BIAS} conditions met the criteria that at least 60 of the combinations observed from a 100-combination sample of the BNC corpus instantiated the most common relation for that head. During our search for predictive heads it emerged that the best candidates were frequently agent nominalizations, and of the 25 heads chosen, 10 were of this type. Combinations were generated so that corresponding stimuli in the P_{BIAS} and $P_{CONTROL}$ conditions had the same modifier and same relation and likewise for the N_{BIAS} and $N_{CONTROL}$ conditions. Each pair of conditions was controlled for word length, plausibility, familiarity and frequency of the head. The average lengths of the P_{BIAS} ($M = 11.3$) and $P_{CONTROL}$ ($M = 11.0$) materials were not significantly different, $t(48) = .39$, $p = .70$. Similarly, there was no reliable difference in the average lengths of the N_{BIAS} ($M = 11.8$) and $N_{CONTROL}$ ($M = 11.3$) materials, $t(48) = .67$, $p = .50$. In a stimulus pre-test two independent judges rated the plausibility of the 100 materials on a scale of 1 to 5, where 5 was the most plausible. The judges were explicitly instructed to only evaluate the concept referred to by the combination and not the manner of its expression. The plausibility ratings did not reliably differ between the P_{BIAS} ($M = 3.4$) and $P_{CONTROL}$ ($M = 3.5$) conditions, $t(48) = -.11$, $p = .22$, or between the N_{BIAS} ($M = 4.0$) and $N_{CONTROL}$ ($M = 4.2$) conditions, $t(48) = -.29$, $p = .77$.

Tagalakakis and Keane (2003) demonstrated that the familiarity of combinations has a large influence on response times in sensibility judgments. Although the familiarity ratings in their experiments were generated by the participants, we reasoned that a corpus study would provide more accurate statistics. At 100 million words, the BNC was not large enough to produce detailed frequencies for novel noun-noun compounds so instead we availed of the internet and gauged combination frequency by using the number of hits generated by a Google search for that combination. Combinations generating less than 50 hits were rejected as being too unusual. One disadvantage to using Google was that many of the hits generated were simply noun co-occurrences and not actual concept combinations. We minimized the number of false positives in the frequency counts by avoiding searches for terms with nouns that commonly double as verbs (e.g. *forest walk*) and also by avoiding searches for regular phrases (e.g. *ocean view* is a popular place name). Using log Google hits, the average frequencies of the P_{BIAS} ($M = 3.6$) and $P_{CONTROL}$ ($M = 3.7$) conditions did not reliably differ, $t(48) = -.34$, $p = .74$. Likewise the average frequencies of the N_{BIAS} ($M = 3.4$)

and $N_{CONTROL}$ ($M = 3.4$) conditions were not reliably different $t(48) = .05$, $p = .96$. Finally, we also controlled for the frequency of the head noun's occurrence as a head by taking the log of the number of combinations in the BNC with the same head. The average head frequencies of the P_{BIAS} ($M = 2.4$) and $P_{CONTROL}$ ($M = 2.2$) conditions were not reliably different, $t(48) = .92$, $p = .36$. Likewise, there was no significant difference between the average head frequencies of the N_{BIAS} ($M = 2.4$) and $N_{CONTROL}$ ($M = 2.2$) conditions, $t(48) = .76$, $p = .45$.

Because our experiment examined the effect of the head's relation type distribution, each combination was categorized using the CARIN taxonomy in order to determine the level of bias of the head towards the instantiated relation. This was not a straightforward task as some of the materials proved difficult to classify, particularly the agent nominalizations. Although Levi (1978) considers these compounds separately from those characterized by recoverably deletable predicates (RDPs), CARIN makes no such distinction. For the sake of comparability with Gagné and Shoben's (1997) experimental method, we endeavored to select the most appropriate category for each combination. However, for materials such as *bridge designer* and *bear hunter*, the final selection did not always satisfactorily reflect the true relationship.

Design A within-participants design was used for the experimental manipulation of condition. Each participant saw the same set of 200 stimuli, comprising the four conditions (P_{BIAS} , $P_{CONTROL}$, N_{BIAS} , $N_{CONTROL}$) of 25 materials each and the 100 nonsensical filler items.

Procedure and Apparatus 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. They were informed that a series of noun-noun compounds would be displayed on the screen for which they would have to make sensibility judgments, pressing J for sense and F for nonsense. Emphasis was placed on the fact that they should only press F if the combination was truly incomprehensible. Each trial was separated by a blank screen lasting for one second. The combination then appeared in the middle of the screen and participants had to make a decision by pressing the appropriate key.

Participants were initially given a short practice session where feedback was given regarding their judgments. The aim of this practice session was to familiarize them with the process of making quick sensibility judgments and also to set a reliable threshold for sensibility. Without such a measure, participants would have been liable to disregard unusual but potentially sensible combinations as nonsense. After completing the practice session, participants were instructed that they were now beginning the experiment. The stimuli were then presented in a random order to each participant.

Results and Discussion

A total of 9.5% of trials were omitted from the analysis. In 6.9% of the trials, the incorrect response was given and hence these trials were not considered. Additionally, responses deemed unreasonably fast ($< 400\text{ms}$, 0.1%) or unreasonably slow (> 4000 , 1.9%) were also excluded. After this initial elimination process, any remaining response times which were more than three standard deviations outside each participant's mean for that condition were also excluded. This excluded another 0.6% of responses.

The mean response times were 1,218; 1,392; 1,370; and 1,348 ms for the P_{BIAS} , P_{CONTROL} , N_{BIAS} and N_{CONTROL} conditions, as illustrated in Figure 2. The mean accuracy rates were .94, .86, .88, and .94 respectively. We conducted several repeated measures ANOVAs in order to examine the influence of the head's relation type distribution on response time, using both participants and items as random factors.

The difference in response times between the P_{BIAS} and P_{CONTROL} conditions was reliable across participants and across items, $F_1(1,35) = 32.30, p < .01$; $F_2(1,24) = 19.02, p < .01$. The difference in accuracy rates between these two conditions was also significant, $F_1(1,35) = 30.83, p < .01$; $F_2(1,24) = 9.04, p < .01$.

The difference between response times in the N_{BIAS} and N_{CONTROL} conditions was not reliable, $F_1(1,35) = .99, p = .33$; $F_2(1,24) = 0.67, p = .42$. However, the difference between the accuracy rates for these conditions was significant, $F_1(1,35) = 21.55, p < .01$; $F_2(1,24) = 4.70, p = .041$.




Figure 2: Response times by condition.

These results demonstrate that the relation type distribution of the head exerts an influence over the ease of interpretation of a combination. Combinations that used a relation that was very typical of the head (e.g. *sausage factory*) were interpreted reliably faster than those that used a head that was less indicative of the instantiated relation (e.g. *sausage machine*). This effect was also reliable when we considered the agent nominalizations and the RDP combinations separately. Conversely, there was no significant difference in response times between combinations where the instantiated relation contradicted the head's bias (e.g. *river factory*) and those involving a neutral head (e.g. *river bench*). Once again, this applied to

both the agent nominalizations and the RDP combinations equally. Taken together, our results suggest that the influence of the head's relation type distribution is limited to situations where it is strongly biased towards the instantiated relation.

Regarding the accuracy rates, we observed significant differences between both sets of conditions. While that between the P_{BIAS} and P_{CONTROL} conditions corresponded with the difference in ease of comprehension, the difference between the N_{BIAS} and N_{CONTROL} conditions requires an alternative explanation given that the response times for these conditions did not differ reliably. We suggest that although the speed of comprehension did not vary, participants may have been more inclined to incorrectly judge the N_{BIAS} stimuli as nonsense. The suggestiveness of the head may have led them to apply the wrong relation and judge the combination on the basis of this incorrect interpretation. For example *river factory* may have initially conjured up the image of a factory that makes rivers. Because the incorrect relation was so apparent, participants may have engaged it ahead of the more sensible yet less obvious alternative. This possibility is evident from the fact that *river factory* produced the lowest accuracy rating of all materials at 36% as opposed to 89% for *river bench*. Despite the discrepancy, both entities have a similar potential to be located next to a river and indeed Google produces 23% fewer hits for the latter phrase. Although these contradictions of the head's bias seem to have increased the error rate by misleading some participants, they did not lead to a significant increase in response time. Hence we conclude that for those materials correctly interpreted, the comprehension process was no more difficult.

Correlation and Regression Analysis In order to determine the correlation between relation strength and response time we applied Gagné and Shoben's (1997) strength equation to the distribution data for the head nouns. This equation, as shown below, consists of an exponential decay function where p_{selected} is the relation frequency of the instantiated relation. While Gagné and Shoben included four terms in the denominator, we included only the proportion of the selected relation, p_{selected} and the relation with the highest remaining proportion, p_1 .

$$\text{Head Strength Ratio} = \frac{e^{ap_{\text{selected}}}}{e^{ap_{\text{selected}}} + e^{ap_1}}$$

The resulting values for head strength encapsulated the information regarding head biases in each of the four conditions. Materials in the P_{BIAS} condition had strengths at the higher end of the range while those in the N_{BIAS} condition were very low because their distributions were dominated by another very frequent relation type. Although Gagné and Shoben (1997) reported that the optimum value for the variable a was about 36, we decided upon a value of

3. Assigning a value of 36 meant that virtually all of the strength values became either a one or a zero, whereas a value of 3 resulted in a more continuous distribution.

We obtained correlations between response time and head strength ($r = -.36, p < .01$), familiarity ($r = -.33, p < .01$), head frequency ($r = -.29, p < .01$), modifier frequency as a modifier in the BNC ($r = -.24, p = .03$), plausibility ($r = -.23, p = .02$) and word length ($r = .10, p = .30$). All the above correlations were significant except that between response time and length.

When considering only the 50 non-biased materials, the correlation between response time and modifier frequency rose to $-.37 (p < .01)$ and that between response time and head frequency rose to $-.38 (p < .01)$. Conversely, the predictive value of strength was greatest for those materials with high strength values. When considering only the 50 materials in the P_{BIAS} and $P_{CONTROL}$ conditions, the correlation between the head strength ratio and response time rose to $-.51 (p < .01)$. For the N_{BIAS} and $N_{CONTROL}$ materials, the correlation between strength and response time was not significant ($r = -.10, p = .48$).

These analyses indicate that past experience of both the head and modifier is important during the comprehension process. The frequency of both constituents correlated significantly with response time, particularly in the absence of a suggestive head bias. However, the degree of correlation between head strength and response time was greatest for combinations with higher strength values. This again suggests that the influence exerted by the head's relation type distribution is greatest for combinations where the instantiated relation is one that is strongly suggested by the head's relation type bias.

We fitted a stepwise regression model using the predictor variables of head strength ratio, familiarity, head frequency, modifier frequency and plausibility as well as the data from all 100 experimental items. The three variables that entered into the model were strength, familiarity and plausibility and the resulting multiple correlation was $.52$. The standardized regression coefficients for the strength, familiarity and plausibility were $-.34, -.27$ and $-.25$ respectively (all $ps < .01$). This supports our hypothesis that the head noun exerts an important influence on those combinations for which it is particularly suggestive.

General Discussion

Our results are consistent with some aspects of the CARIN model, in so far as they indicate that interpretation is influenced by relation type distribution. The results support CARIN's principle that it is easier to determine whether a phrase makes sense if it involves a relation that is highly available. However, several refinements to the theory are required in order to account for the role of the head. Our results contradict CARIN's assertion that the availability of a relation for the head noun has little or no effect on interpretation, hence undermining the notion of modifier primacy. The finding that both constituents in a combination have equal potential to influence comprehension challenges

the idea that modifiers have certain semantic privileges in determining the relation.

Lack of Prior Evidence

No previous studies have revealed a facilitating influence exerted by the head's relation type distribution. We believe that this may have been due in part to the experimental materials used and in part to the relation distributions derived using these materials. Whereas we purposely selected the most biased heads for our experiment, Gagné and Shoben's (1997) method was far less selective in that it involved an arbitrary set of heads that combined using their highest frequency relations. As a result, the majority of these heads are not associated with highly dominant relations and taking BNC frequencies into account, only four of them actually meet our criteria for sufficient bias (*magazine, equipment, crisis* and *scandal*). If our hypothesis that only the most predictive heads exert an influence is accurate, then this may have contributed to the lack of an effect.

Another reason for the discrepancy may be due to the fact that the influence of the head was not assessed independently. Conditions of high modifier strength were combined with conditions of low head strength and vice versa, with the influence of one condition potentially obscuring that of the other. Furthermore, we have reason to believe that predictive heads may be far rarer than predictive modifiers, an inequality that would consequently be exacerbated by a random sample of materials. Our reasoning is as follows. The concept denoted by a particular sense of a head noun is of some ontological category which is also therefore the ontological category of the entire combination. This category is generally open to modification along many dimensions as the concept has many features with default values. While a modifier noun also denotes a concept of some ontological category, that concept can serve as a fitting value for only a limited number of modifying dimensions and will typically modify the same dimension in a wide variety of head nouns.

For example, if we consider a noun such as *dog*, any aspect of the dog concept can be potentially modified, including where the dog lives, what it eats, what it looks like, what it smells like etc. However, when *dog* is used as a modifier, it can be employed in a more limited number of ways. Its usage signifies that a separate concept is related to dogs in some way, and the possible variations on this relation are fewer. For most nouns, the number of dimensions in which they can be modified will be considerably greater than the number of dimensions for which they can serve as modifiers. Additionally, a wide range of head nouns will be modified in a similar way by the same modifier. As a result, it will be relatively easy to find modifying nouns with skewed relation type distributions but harder to find head nouns with skewed distributions. In a random sample, the few head nouns that might markedly reveal such a tendency are likely to be

outnumbered by others that will not, hence diminishing the effect observable by experiment.

This phenomenon has repercussions for relation type distributions that are derived from arbitrary pairings as it implies that such a technique will be less successful at emulating the head's natural combinatorial disposition than that of the modifier. The greater flexibility of head nouns means they have a greater potential to be modified in ways that are simultaneously sensible yet unnatural. As a result, predictive heads will often need to be paired with a precise domain of modifiers in order to reveal their bias. Predictive modifiers, on the other hand, are more likely to betray their bias in arbitrary pairings because of their ability to modify the same slot in a wide variety of head nouns, regardless of the nature of the head concept.

As an example let us consider one of our predictive head materials, *soup*. The BNC reveals that 93% of the time, this head combined with a modifier referring to what the soup contained. However, if we look at the pairings of *soup* with Gagné and Shoben's 91 modifiers we notice that in order for the <has> relation to emerge, *soup* must be paired with a modifier that denotes a food substance. Since only 12 of Gagné and Shoben's modifiers fall into this category, *soup* gets a far lower weighting for the <has> relation than we would expect. Additionally, when we consider the other sensible combinations that emerge, we see that other attributes of the *soup* concept are affected, resulting in modifications that are unrepresentative of its typical usage as a head. Thus we find combinations such as *party soup*, *home soup*, *office soup*, *family soup*, and *student soup*. These uncharacteristic modifications give *soup* a misleadingly high weighting for the <located> and <for> relations. The use of arbitrary pairings means that the emergent distributions are strongly influenced by relation type biases among the set of modifiers and hence the only biased heads that can be reliably detected using this technique are those that have attributes with particularly versatile domains like *book* or *magazine*.

Modifier Influence

While we uncovered a non-linear relationship between the degree of head influence and head strength, it is unclear whether the same pattern applies to modifier influence. Gagné and Shoben (1997) dichotomized relation frequencies as high or low and as a result, information about differences in frequencies was lost, making it difficult to tell whether the influence of the modifier extends linearly across all modifier strength values. We propose that only the relation distributions of extremely biased constituents feature in the comprehension process, since less evident distributional variations would be of little predictive value. Consequently, while the CARIN model assumes that the relation distribution of every relation is stored separately for every single noun, we propose that people encode only that information which is likely to be of benefit during interpretation, specifically those distributions that are particularly biased. Keeping separate distributions for every

noun in the lexicon would be excessively inefficient and, given our finding that only heads with the greatest bias influenced response time, it seems plausible that speakers might only be aware of the clearest distinctions in a noun's relation type distribution.

Conclusion

Experiment has shown that the ease of comprehension of a conceptual combination is influenced by the relation type distribution of heads that are particularly suggestive of the instantiated relation. This finding contradicts one of the central tenets of the CARIN theory, challenging the notion that modifiers enjoy certain semantic privileges in determining the relation (Gagné & Shoben, 1997). On the other hand, we found no influence on comprehension when the instantiated relation contradicted the head's bias and correlation analyses between response time and the head strength ratio were only significant for the positively biased condition. This suggests that the foremost influence of the head's relation type distribution is to speed up interpretations under the presence of an extremely dominant relation.

While our study has shown that the head noun can influence ease of comprehension, it is possible that in general, modifiers are better predictors of the relation. The discrepancy between the modifier's modifying potential and the head's potential to be modified means that biased modifiers are far more common than biased heads, and the greater predictive value of the modifier in general may explain why Gagné and Shoben's (1997) random sample of heads and modifiers revealed an influence for the modifier's relation type distribution but not for the head. This asymmetry in the predictive potential of both constituents is an interesting phenomenon that may reveal much about the semantic structure of nominal concepts and their representation in the lexicon.

References

- Corley, S., Corley, M., Keller, F., Crocker, M. W. & Trewin, S. (2001). Finding syntactic structure in unparsed corpora: The Gsearch corpus query system. *Computers and the Humanities*, 35, 81-94.
- Gagné, C. L. & Shoben, E. J. (1997). Influence of thematic relations on the comprehension of modifier-noun combinations. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 23, 71-87.
- Maguire, P. & Cater, A. (2004). Is conceptual combination influenced by word order? *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics*
- Storms, G. & Wisniewski, E. J. (in press) Does the order of head noun and modifier explain response times in conceptual combination? *Memory and Cognition*.
- Tagalakis, G. & Keane, M. T. (2003). Modeling the understanding of noun-noun compounds: The role of familiarity. *Proceedings of the European Cognitive Science Conference*.