A Concept-Centered Approach to Noun-Compound Interpretation

Cristina Butnariu

School of Computer Science and Informatics
University College Dublin
Belfield, Dublin 4
Ioana.Butnariu@ucd.ie

Tony Veale

School of Computer Science and Informatics
University College Dublin
Belfield, Dublin 4
Tony.Veale@ucd.ie

Abstract

A noun-compound is a compressed proposition that requires an audience to recover the implicit relationship between two concepts that are expressed as nouns. Listeners recover this relationship by considering the most typical relations afforded by each concept. These relational possibilities are evident at a linguistic level in the syntagmatic patterns that connect nouns to the verbal actions that act upon, or are facilitated by, these nouns. We present a model of noun-compound interpretation that first learns the relational possibilities for individual nouns from corpora, and which then uses these to hypothesize about the most likely relationship that underpins a noun compound.

1 Introduction

Noun compounds hide a remarkable depth of conceptual machinery behind a simple syntactic form, *Noun-Noun*, and thus pose a considerable problem for the computational processing of language (Johnston and Busa, 1996). It is not just that compounds are commonplace in language, or that their interpretation requires a synthesis of lexical, semantic, and pragmatic information sources (Finin, 1980); compounds provide a highly compressed picture of the workings of concept combination, so there are as many ways of interpreting a noun compound as there are ways of combining the underlying concepts (Gagné, 2002). Linguists have thus attempted to

understand noun-compounds as full propositions in which a phrase with two nouns connected by an explicit relation – usually expressed as a verb and a preposition – is compressed into a pair of nouns (Levi, 1978). Since these noun-pairs must allow an audience to reconstruct the decompressed proposition, there must be some systematic means by which the missing relation can easily be inferred.

This framing of the problem as a search for a missing relation suggests two broad strategies for the interpretation of compounds. In the first, the top-down strategy, we assume that there are only so many ways of combining two concepts; by enumerating these ways, we can view the problem of interpretation as a problem of classification, in which compounds are placed into separate classes that each correspond to a single manner of concept connection (Kim and Baldwin, 2006), (Nastase and Szpakowicz, 2003). This strategy explicitly shaped the SemEval task on classifying semantic relations between nominals (Girju *et al.*, 2007) and so is employed by all of the systems that participated in that task. In the second, the bottom-up strategy, we assume that it is futile to try and enumerate the many ways in which concepts can relationally combine, but look instead to large corpora to discover the ways in which different word combinations are explicitly framed by language (Nakov, 2006), (Turney, 2006a).

In this paper we describe an approach that employs the bottom-up strategy with an open-rather than closed-inventory of inter-concept relations. These relations are acquired from the analysis of large corpora, such as the Web IT corpus of Google n-grams (Brants and Franz, 2006). We argue that an understanding of noun-compounds requires an understanding of lexicalized concept combination, which in turn requires an understanding of how lexical concepts can be used and connected to others. As such, we do not use corpora as a means of

characterizing noun-compounds themselves, but as a means of characterizing the action possibilities of the individual nouns that can participate in a compound. In other words, we attempt to characterize those properties of different concepts denoted by nouns to help predict how those nouns will combine with others, and through which relations. For instance, “diamonds” can be used to cover and encrust jewelry or to provide a sharp tip for various tools; we see the former usage in “diamond bracelet” and the latter in “diamond saw”. Likewise, “cheese” is a solid substance which can be cut, as in “cheese knife”, an edible substance that can be used as a filling, as in “cheese sandwich”, and a substance that can be melted as a topping, as in “cheese pizza”. It follows that a sandwich can be filled, a pizza can be topped, knives can cut and a bracelet can have a covering of gems. We use relational possibilities as a general term for what are sometimes called the qualia of a word (Pustejovsky, 1995), and learn the linguistic relational possibilities of nouns by seeking out specific textual patterns in corpora.

In section 2 we consider the most currently relevant elements of the substantial body of past work in this area. In section 3 we describe how corpus analysis is used to identify the most common lexico-semantic relational possibilities of nouns, while in section 4 we describe how these relations are used, in conjunction with web-based validation, to interpret specific noun-compounds. We present an evaluation of this approach in section 5 and conclude the paper with some final remarks in section 6.

2 Related Work

Machine-learning and example-based approaches to noun-compounds generally favor the top-down strategy for defining relations, since it allows training data and exemplars/cases to be labeled using a fixed inventory of relational classes. As noted earlier, this strategy is characteristic of the systems that participated in the SemEval task on classifying semantic relations between nominals (Girju *et al.*, 2007), such as Butnariu and Veale (2007). Though the inventory is fixed in size, it can be defined using varying levels of abstraction; for instance, Nastase and Szpakowicz (2003) use an inventory of 35 relations, 5 of which are top level relations with the remaining 30 at the lower level. The top-down strategy pre-dates these computational approaches, and is a key aspect of the

foundational work of Levi (1978) and of subsequent work by Gagné and Shoben (1997), both of whom posit a small set of semantic relations as underpinning all noun compounds. More recently, Kim and Baldwin (2005) use a fixed inventory of semantic relations to annotate a case-base of examples; new compounds are understood by determining their lexical similarity to the closest annotated compounds in the case-base. Kim and Baldwin (2006) link their relations to specific seed verbs that linguistically convey these relations, and then train a classifier to recognize which semantic relation is implied by a pair of nouns connected by a given intervening verb. This approach appears to be sensitive to the number of seed verbs; on a test involving 453 noun compounds, an accuracy of 52.6% is achieved with 84 seed verbs, but just 46.6% with 57 seed verbs.

Verbs understandably play a key role in the interpretation of compounds, since some kind of predicate must be recovered to link both nouns. For instance, Levi (1978) uses verbs to make explicit the implicit relation between the nouns of a compound, while Finin (1980) characterizes the relation in a noun-noun compound using an inventory of all the possible verbs that can link both nouns; thus, e.g. *salt water* is interpreted using a relation like *dissolved in*. Nakov (2006) takes a similar approach, and uses verb-centred paraphrases to express the semantic relations between the nouns of a compound. He argues that the meaning of a compound is best expressed via a collection of appropriate verbs rather than via the abstract relations (such as **Cause**, **Location**) that are used in more traditional approaches, such as those of Levi (1978) and Gagné (2002).

Nakov (2006) pursues a bottom-up strategy in which an open-ended inventory of relations is discovered using linguistic evidence. Turney (2006a, 2006b) similarly pursues a bottom-up, data-driven approach, in which semantic relations are expressed via representative lexico-syntactic patterns that are mined from large text corpora. Turney (2006a) sorts these relational patterns by *pertinence*, a measure that reflects the similarity of the noun pairs in the corpus in which each pattern is observed to occur. Patterns which are relatively unambiguous and which serve to cluster noun pairs with similar meanings have higher pertinence than those that do not.

The approach described here is similarly corpus-based and verb-centric, but it is also noun-centric rather than pair-centric, which is to say, we use

corpus analysis to learn about the relational behavior of individual nouns rather than pairs of nouns. Like many other authors, from Finin (1980) to Nakov (2006), we see the problem of compound interpretation as a problem of paraphrase generation, in which a suitable verb (with an optional preposition) is used to linguistically re-frame the compound as a complete proposition. This linguistic frame is a relational possibilities of one of the nouns that is apt for the other. Following Gagné and Shoben (1997), this relational possibilities is frequently suggested by the modifier noun, but as we now describe, it may also be suggested by the head.

3 Acquisition of Relational Possibilities

The meaning of a noun compound can be paraphrased in a variety of ways. For instance, consider the compound “headache pill”, which might be paraphrased as follows:

- P1: headache-inducing pill
- P2: headache prevention pill
- P3: pill for treating headaches
- P4: pill that causes headaches
- P5: pill that is prescribed for headaches
- P6: pill that prevents headaches

Some paraphrases are syntactic variants of others (e.g., P2 and P6), others employ lexical variation (e.g., P1 and P4) and others are co-descriptions of the same event (e.g., P3 and P5 or P5 and P6). It thus seems unreasonable to try and reduce these meanings to a single semantic relation, since the compound can be used to mean several of P1 ... P6 simultaneously. Rather than try to construct an inventory of logical relations, closed or otherwise, we shall instead treat linguistic frames like “for treating X”, “that prevents X”. etc. as proxies for the relations themselves, while retaining the capacity to treat syntactic variants as proxies for the same relation. Moreover, these linguistic frames are relational possibilities of specific words, so that “-inducing X” is a relational possibility of “headache” while “for treating X” is a relational possibility of “pill”. Thus, a compound of the form “headache X” might be re-framed as “headache-inducing X” and a compound of the form “X pill” might be re-framed as “pill that prevents X”, “pill that causes X” or “pill for treating X”.

The relational possibilities of individual words can be acquired from a large n-gram corpus like that of Brants and Franz (2006), as derived from

Google’s web index. Table 1 summarizes the linguistic relational possibilities that can be derived from specific n-gram patterns.

<i>Google n-gram pattern</i>	<i>Relational possibilities</i>	<i>Logical Form</i>
X – Verb+ing	X Verb+ing Y	verb(X, Y)
X – Verb+ed	X Verb+ed Y	verb(X, Y)
Verb+ed prep X	Y Verb+ed prep X	verb_prep(Y, X)
X Verb+ed	X Verb+ed prep Y	verb_prep(X, Y)
for Verb+ing X	Y for Verb+ing X	verb(Y, X)
X for Verb+ing	X for Verb+ing Y	verb(X, Y)
that Verb+s X	Y that Verb+s X	verb(Y, X)
X that Verb+s	X that Verb+s Y	verb(X, Y)

Table 1. For an anchor noun *X*, the n-gram (left) suggests relational possibilities (middle) to link to a generic noun *Y*; different linguistic relational possibilities can have the same logical form (right).

For example, we extract the following linguistic relational possibilities for the noun “diamond”, where Google frequencies are given in parentheses:

```

accented_with_diamonds(4224),
encrusted_with_diamonds(3990),
decorated_with_diamonds(2616),
  based_on_diamond(2148),
covered_with_diamonds(2018),
filled_with_diamonds(1942),
adorned_with_diamonds(1462),
coated_with_diamond(1150),
for_buying_diamonds(618),
for_grading_diamonds(342),
for_cutting_diamonds(430),
dusted_in_diamond(168),
bedecked_with_diamonds(140),
tipped_with_diamond(108),
crowned_with_diamonds(98),
for_exporting_diamonds(98),
embossed_with_diamond(90),
edged_with_diamonds(82),
drilled_with_diamond(86),
that_sells_diamonds(44)...
```

A hat may be *crowned* with diamonds, a watch *decorated* with diamonds, a bracelet *covered* with diamonds, a throne encrusted with diamonds and a king *bedecked* with diamonds – each is an elaboration of a basic *covering* relation, but each adds nuances of its own that we do not want to lose in an interpretation that is

maximally specific to the nouns concerned. We therefore take the view that relations should be as open-ended and nuanced as the linguistic evidence suggests; if one needs to see two different relations as broadly equivalent, resources like WordNet (Fellbaum, 1998) can be used to make the generalization.

4 Interpreting Noun Compounds

We see interpretation of a compound M-H as a two-stage process of divergent generation followed by convergent validation. The generation process simply considers the relational possibilities associated either with the modifier M or the head H and generates a paraphrase from each. Consider the compound “yeast bread” where M = “yeast” and H = “bread”; the relational possibilities for “yeast” and “bread” are used to generate a set of potential paraphrases as shown in Table 2. For clarity, “M” and “H” denote the parts of each paraphrase frame that will be filled with the modifier and head respectively.

Relational possibilities for M	Paraphrases for M-H
H Verb+ed prep M <i>e.g., H derived from yeast</i>	H Verb+ed prep M <i>e.g., bread derived from yeast</i>
H that Verb+s M <i>e.g., H that contains yeast</i>	H that Verb+s M <i>e.g., bread that contains yeast</i>
Relational possibilities for H	Paraphrase for M-H
H Verb+ed prep M <i>e.g., bread prepared with M</i>	H Verb+ed prep M <i>e.g., bread prepared with yeast</i>
H that Verb+s M <i>e.g., bread that has M</i>	H that Verb+s M <i>e.g., bread that has yeast</i>

Table 2. Relational possibilities of the head (H) and modifier (M) nouns used for paraphrasing.

The relational possibilities for the head noun “bread” yield the following paraphrases, where numbers in parentheses are Google frequencies for the original n-grams on which each paraphrase is based:

“bread made from yeast” (6335), “bread topped with yeast” (6043), “bread made with yeast” (4726), “bread stuffed with yeast” (3871), “bread baked in yeast” (3341), bread made of yeast” (3064), “bread served

with yeast” (3012), “bread soaked in yeast” (2975), “bread dipped in yeast” (2873), “bread filled with yeast” (2783), ...

Similarly, the relational possibilities for the modifier noun “yeast” yield the following paraphrases:

“bread expressed in yeast” (14058), “bread leavened with yeast” (10816), “bread derived from yeast” (2562), “bread based on yeast” (1200), “bread fermented with yeast” (842), “bread raised with yeast” (736), “bread induced in yeast” (342), “bread infected with yeast” (262), “bread filled with yeast” (120), ...

While these two sets of relational possibilities capture the most salient activities in which “yeast” and “bread” participate, many of the paraphrases listed here are inappropriate for “yeast bread”. Candidate paraphrases for a noun compound are useful only when one has a means of determining the degree to which paraphrases are meaningful and apt and of rejecting those which are not. This process typically assumes that a meaningful paraphrase is one for which evidence of prior usage can be found in a large corpus (like the web); the greater this evidence, the more favored a given paraphrase should be. This assumption is central to Nakov (2006), who uses templates to find paraphrases for a noun compound on the web. These templates use the Google wildcard * to indicate the position of a verb so that the specific verbs at the heart of a paraphrase can be mined from the snippets that are returned. Nakov (2007) uses the schematic patterns “N1 that * N2”, “N2 that * N1”, “N1 * N2” and “N2 * N1”, where the wildcard can stand for as many as eight contiguous words.

Relational possibilities allow us, in the first divergent stage of interpretation, to generate fully-formed paraphrases that do not require wildcards, so the second convergent stage of interpretation simply needs to validate these paraphrases by finding one or more instances of each on the web. Indeed, an especially compelling paraphrase may be found in the Google n-grams themselves, without recourse to the web. For instance, the paraphrase “bread leavened with yeast” has a frequency of 56 in the database of Google 4-grams, while the

paraphrase “bread based on yeast” has such a low web frequency of 2 hits that it can be validated only by going to the web.

But web-based validation has its limitations: it cannot account for novel and creative compounds, nor can it account for conventional compounds whose meaning is not echoed in an expanded paraphrase-form on the web. Thus, we also consider an alternate validation procedure for those paraphrases that can be generated both from a modifier noun relational possibility *and* from a head noun relational possibility. For example, “bread filled with yeast” can be derived from the head relational possibility “bread filled with X” which has a frequency of 2783, and from the modifier relational possibility “X filled with yeast” which has a frequency of just 120. This dual basis for generation provides evidence that the paraphrase is meaningful without the need to actually find the paraphrase on the web. We refer to the validation of paraphrases in this way as validation by *matching relational possibilities* of the modifier and head nouns.

This *matching relational possibilities* procedure does not require web validation, and so does not produce a web frequency for each paraphrase. We thus need to assign a score to a paraphrase based on the web frequencies of the matching relational possibilities that give rise to it. For simplicity, we add the web frequency of the head relational possibility (e.g., 2783 from “bread filled with X”) to the frequency of the modifier relational possibility (e.g., 120 from “X filled with yeast”) to obtain an invented frequency for the generated paraphrase (e.g., 2903 for “bread filled with yeast”).

The third and more restricted validation procedure we employ is a hybrid one, based on the intersection of the two procedures above: we require web-validation of paraphrases that are already validated by virtue of arising from matching head and modifier relational possibilities. In this case, we rank the paraphrases by their actual web frequency. The set of paraphrases validated by the hybrid approach will be a subset of the paraphrases validated by the other two validation methods; the size of this subset will be informative about the relative utility of each procedure.

5 Empirical Evaluation

To evaluate the relational possibility approach to noun-noun interpretation, we perform two experiments: one to consider how well the set of

validated paraphrases can be mapped to the abstract relations used by (Nastase and Szpakowicz, 2003) to annotate their noun-noun compounds, and one to consider how well these paraphrases match the paraphrases offered by humans for the same noun compounds. To understand the role of different validation strategies, we use three variants of the model that correspond to the three means of validating paraphrases: model-1 uses the presence of the relational possibility on the web as the mark of a valid paraphrase; model-2 uses the matching relational possibilities procedure to validate a paraphrase (i.e., the paraphrase must arise from both an relational possibility of the modifier *and* of the head); a third model, model-3 intersects both validation procedures. In each case, validated paraphrases are ranked by their frequency scores, as found explicitly on the web in the case of model-1 and model-3, or as invented in model-2.

5.1 Mapping compounds to abstract relations

In the first experiment, we test the relational possibility model on a set of noun-noun compounds from Nastase and Szpakowicz (2003), whose data is pre-classified into abstract classes of semantic relations (i.e., **Agent**, **Instrument**, **Location**). We perform a manual analysis on the paraphrases that are generated and validated for each noun pair, to measure how accurately each paraphrase matches the pre-classified abstract semantic relation. The Nastase and Szpakowicz (2003) dataset comprises 600 word pairs of the form adj-noun, adv-noun and noun-noun; for this experiment we use only the 329 noun-noun pairs, which are each pre-labeled with one of 28 different semantic relations.

We consider and quantify two eventualities here: those situations in which the relational possibility model generates and validates a paraphrase that closely corresponds to the semantic relation assigned by Nastase and Szpakowicz (2003); and those situations in which the relational possibility model generates and validates an interpretation that a human judge considers a plausible and sensible interpretation of a compound regardless of Nastase and Szpakowicz (2003)’s interpretation. Table 3 presents validated relational possibilities for the compound “olive oil”, where those that match the pre-classified relation are in bold, and those that are otherwise plausible are italicized.

<i>Paraphrases generated by web-based validation</i>	<i>Paraphrases generated by matching relational possibilities</i>	<i>Paraphrases generated by Nakov (2007)</i>
extracted from (189), obtained from (132), <i>mixed with</i> (87), made from (75), produced from (38), pressed from (35), <i>colored</i> (25), <i>infused</i> (20), <i>enriched with</i> (16), made of (14), <i>flavored</i> (13), made with (12), derived (10), based (10), <i>produced by</i> (9), <i>blended with</i> (8), <i>coloured</i> (7), based on (7), <i>combined with</i> (6), found in (6), <i>dissolved in</i> (6), <i>served with</i> (6), contained in (5), <i>replaced by</i> (4), <i>flavoured</i> (3), come from (3)...	used in (25839), obtained from (15352), extracted from (14561), made from (11627), found in (11524), <i>used for</i> (9919), <i>mixed with</i> (9781), produced from (7794), <i>produced by</i> (6776), made with (5423), <i>used as</i> (4880), are in (4577), contained in (4551), come from (4241), based on (4135), <i>combined with</i> (4029), <i>added to</i> (3848), <i>made in</i> (3608) ...	come from (13), be obtained from (11), be extracted from (10), be made from (9), be produced from (7), be released from (4), <i>taste like</i> (4), be beaten from (3), be produced with (3), emerge from (3)

Table 3. Validated paraphrases for “olive oil”; matches with Nastase and Szpakowicz are in bold; other sensible interpretations are italicized.

We also consider the rank of the paraphrases that match the relations assigned by Nastase and Szpakowicz (2003) to this data set. Figure 1 graphs the F-measure for the relational possibilities approach when this relation is the top-ranked validated paraphrase, when it is in the top two validated paraphrases, and more generally, when it is in the top n validated paraphrases, $n \leq 20$. Model-1 (web-based validation) out-performs Model-2 (matching relational possibilities, with no web validation) when we consider just a small window of top ranked paraphrases, but this situation reverses as the window (whose size is given on the x-axis) is enlarged.

Figure 1. F-measure for target semantic relations of top n ranked paraphrases generated with Model-1, Model-2 and Model-3.

Model 3 (which requires both matching relational possibilities and web validation) shows similar results to Model 2 (web validation only), which suggests that the matching relational possibilities criterion is strongly predictive of web-validation. This further suggests that matching relational possibilities alone can reliably validate a paraphrase even when web

evidence is lacking, as will be the case in creative noun compounds.

During this evaluation process, we observe a tendency for specific paraphrases to co-occur when conveying a certain relation. For instance, *Y obtained from X* typically co-occurs with *Y produced from X* to indicate Nastase and Szpakowicz’s **Source** relation, while *Y caused by X* co-occurs with *Y induced by X* to convey their **Effect** relation, and *Y owned by X* co-occur with *Y held by X* to indicate their **Possessor** relation. This observation is similar to that of Nakov (2007), who performs a manual analysis of paraphrases obtained from web-mining. The results he reports are similar to those obtained using the relational possibilities approach, as shown in Table 3.

5.2 Comparing human-generated paraphrases

In the second experiment, we compared the paraphrases validated by the relational possibilities approach to human-generated paraphrases reported by Nakov (2007) and to the paraphrases generated by Nakov’s own web-mining approach to this task. Nakov (2007) collected human paraphrases for each noun-compound in his data-set (250 noun compounds listed in the appendix of Levi, 1978) by asking subjects to rephrase a noun-compound using a relative-clause centred around a single verb with an optional preposition. This rephrasing elicited human-generated paraphrases like the following:

'neck vein is a vein that comes from the neck'
'neck vein is a vein that drains the neck'

Nakov then extracted normalized verbs and prepositions from these paraphrases to obtain a reduced verb-based form for each, e.g., to obtain the reduced forms *come from* and *drain* from the above examples. He used 174 subjects for this task, to generate around 17,000 reduced forms, or 71 forms per compound.

For each of his 250 noun pairs we constructed three vectors h , w , and a , using human-generated paraphrase verbs and their frequencies (h), Nakov's web-extracted verbs and their frequencies (w) and the verbs of the paraphrases obtained using the relational possibilities approach and their frequencies (a). Following Nakov (2007), we then calculated the cosine correlation between two frequency vectors using the formula:

$$sim_{cos}(h,w) = \sum h_i w_i / \sqrt{\sum h_i^2} \sqrt{\sum w_i^2}$$

For ease of comparison, the a vector is populated with verbs and frequencies from just two patterns, $Y \text{ Verb}+ed \text{ Prep } X$ and $Y \text{ that Verb}+s X$. In Table 4 we report the average cosine correlation across the vectors for all 250 noun pairs, to compare for the three validation models the relational possibilities-based and Nakov's web-generated paraphrases and the relational possibilities -based and human-elicited paraphrases. Also shown, in the last row, is the average cosine correlation between Nakov's web-mined paraphrases and human-elicited paraphrases, as reported in Nakov (2007).

Model 1 (web-validation)	
<i>correlation to humans</i>	26.8 %
<i>correlation to web-mined approach</i>	27.1 %
Model 2 (matching relational possibilities)	
<i>correlation to humans</i>	17 %
<i>correlation to web-mined approach</i>	14.25 %
Model 3 (intersection of Model-1 and Model-2)	
<i>correlation to humans</i>	27.9 %
<i>correlation to web-mined approach</i>	28 %
Web-mining (Nakov, 2007)	
<i>correlation to humans</i>	31.8%

Table 4. Average correlation between web-mined paraphrases and relational possibilities-based paraphrases with human elicitation.

The results show the difference in quality of the paraphrases validated by each of our models. The matching-relational possibilities model

(Model 2) yields the largest number of paraphrases. In the first experiment we showed that this model outperforms the other two when we consider just top-ranked paraphrases, but here it appears that this wider range of potentially creative interpretations diminishes the cosine correlation with human-elicited interpretations. But the most plausible paraphrases come to the fore in the hybrid model (Model 3), whose paraphrases are a subset of those of Models 1 and 2. This hybrid approach also outperforms Model 1 and compares well with the results obtained by web mining.

The difference in cosine correlation between human-elicited and relational possibilities-based paraphrases in Model-3 (27.9%) and Nakov's web-mined and human-elicited paraphrases (31.8%) can be justified both by the type of patterns used in the comparison, and by the type of patterns used to validate paraphrases. For one, we consider paraphrases generated using just two forms of relational possibilities, $Y \text{ Verb}+ed \text{ Prep } X$ and $Y \text{ that Verb}+s X$, since these can be directly compared to the type of relative-clause paraphrases used in this experiment. Furthermore, relational possibilities are derived from Google n-grams where $n < 6$, we allow up to four words to intervene between the modifier and the head in a paraphrase, while the web-mining paraphrases benefit from a larger window of intervening words (up to 8). Nonetheless, in 88 out of the 250 pairs, the correlation between relational possibilities-based and human-elicited paraphrases is larger than that observed for the web-mining approach.

6 Conclusions

Since the meaning of noun compounds arises from a combination of individual noun meanings, it follows that the key input to the process of compound interpretation is detailed linguistic knowledge about how these nouns are conventionally used in language. This point may seem obvious, but a model of compounding can place so much emphasis on the behavior of noun-pairs that the linguistic behavior of nouns in isolation is easily over-looked.

We have presented a model of noun-compounding that places nouns and their specific linguistic relational possibilities at the centre of processing. When one considers that linguistic relational possibilities capture aspects of noun meaning such as purpose, constitution and agency, their realization here can be viewed as a

generalized and lexicalized aspect of qualia structure in the sense of Pustejovsky (1995) and Johnston and Busa (1996). Indeed, the n-gram patterns used to extract these relational possibilities from corpora are not unlike the patterns used by Cimiano and Wenderoth (2007) to harvest qualia structures from the web.

We conclude from the empirical observation that the hybrid model outperforms the web-based model (albeit slimly) in experiment 2 while both perform equally well in experiment 1, is that the modifier and head are of comparable performance when paraphrasing the interpretations of noun compounds. Recall that the web-validation approach (Model-1) generates interpretations from either the modifier or the head, while the matching-relational possibilities and hybrid models require both to contribute equally.

Necessary extensions to the approach include the acquisition of more relational possibilities of greater linguistic complexity, the ability to organize relational possibilities hierarchically according to their underlying semantic meanings, and the ability to recognize an implication structure among different but related relational possibilities.

Acknowledgement

We would like to thank Preslav Nakov for providing us the data used in the second experiment.

References

- Brants, T., and Franz, A. 2006. *Web 1t 5-gram version 1*, Linguistic Data Consortium.
- Butnariu, C., and Veale T. 2007. A hybrid model for detecting semantic relations between noun pairs in text. In *Proc. of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, Prague., Association for Computational Linguistics.
- Cimiano, P., and Wenderoth, J. 2007. Automatic Acquisition of Ranked Qualia Structures from the Web. In *Proc. of the 45th Annual Meeting of the ACL*, pp 888-895.
- Fellbaum, C. 1998. *WordNet, an electronic lexical database*. Cambridge: MIT Press.
- Finin, T. 1980. *The semantic interpretation of compound nominals*. Urbana, Illinois: University of Illinois dissertation.
- Gagné, C. L., and 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.
- Gagné, C. L. 2002. Lexical and Relational Influences on the Processing of Novel Compounds. *Brain and Language* 81(1-3), pp 723-735.
- Girju, R., Nakov, P. Nastase, V., Szpakowicz, S., Turney, P., and Yuret, D. 2007. Semeval-2007 task 04: Classification of semantic relations between nominals. In *Proc. of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, 13–18, Prague, Czech Republic. ACL.
- Johnston, M., and Busa, F. 1996. Qualia structure and the compositional interpretation of compounds. In *Proc. of the ACL SIGLEX workshop on breadth and depth of semantic lexicons*, Santa Cruz, CA.
- Kim, S. N., and Baldwin, T. 2006. Interpreting semantic relations in noun compounds via verb semantics. In *Proc. of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the ACL*, pp 491–498, NJ, USA.
- Kim, S. N., and Baldwin, T. 2005. Automatic interpretation of noun compounds using WordNet similarity. In *Proc. of the 2nd International Joint Conference On Natural Language Processing*, pp 945–956, Cheju, Korea.
- Levi, J. 1978. *The syntax and semantics of complex nominals*. NY: Academic Press.
- Nakov, P., and Hearst, M. A. 2006. Using verbs to characterize noun-noun relations. In *AIMSA, Jerome Euzenat and John Domingue (eds.), vol. 4183 of Lecture Notes in Computer Science*, pp 233–244. Springer.
- Nakov, P. 2007. *Using the Web as an Implicit Training Set: Application to Noun Compound Syntax and Semantics*. Ph.D. Dissertation, University of California at Berkeley.
- Nastase, V., and Szpakowicz, S. 2003. Exploring noun-modifier semantic relations. In *Proc. of the 5th International Workshop on Computational Semantics (IWCS-5)*, pp 285–301, Tilburg, The Netherlands.
- Pustejovsky, J. 1995. *The Generative Lexicon*. The MIT Press, Cambridge, MA.
- Turney, P. 2006a. Expressing implicit semantic relations without supervision. In *Proc. of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the ACL*, pp 313–320, NJ, USA.
- Turney, P. D. 2006b. Similarity of semantic relations. *Computational Linguistics*, 32, pp 379–416.