Making Lexical Ontologies Functional and Context-Sensitive

Tony Veale
Computer Science and Informatics
University College Dublin
Ireland
tony.veale@ucd.ie

Yanfen Hao
Computer Science and Informatics
University College Dublin
Ireland
yanfen.hao@ucd.ie

Abstract

Human categorization is neither a binary nor a context-free process. Rather, some concepts are better examples of a category than others, while the criteria for category membership may be satisfied to different degrees by different concepts in different contexts. In light of these empirical facts, WordNet's static category structure appears both excessively rigid and unduly fragile for processing real texts. In this paper we describe a syntagmatic, corpus-based approach to redefining WordNet's categories in a functional, gradable and context-sensitive fashion. We describe how the diagnostic properties for these definitions are automatically acquired from the web, and how the increased flexibility in categorization that arises from these redefinitions offers a robust account of metaphor comprehension in the mold of Glucksberg's (2001) theory of category-inclusion. Furthermore, we demonstrate how this competence with figurative categorization can effectively be governed by automatically-generated ontological constraints, also acquired from the web.

1 Introduction

Linguistic variation across contexts is often symptomatic of ontological differences between contexts. These observable variations can serve as valuable clues not just to the specific senses of words in context (e.g., see Pustejovsky, Hanks and Rumshisky, 2004) but to the underlying ontological structure itself (see Cimiano, Hotho and Staab, 2005). The most revealing variations are syntagmatic in nature, which is to say, they look beyond individual word forms to larger patterns of contiguous usage (Hanks, 2004). In most contexts, the similarity between chocolate, say, and a narcotic like heroin will meagerly reflect the simple ontological fact that both are kinds of substances; certainly, taxonomic measures of similarity as discussed in Budanitsky and Hirst (2006) will capture little more than this commonality. However, in a context in which the addictive properties of chocolate are very salient (e.g., an online dieting forum), chocolate is more likely to be categorized as a drug and thus be considered more similar to heroin. Look, for instance, at the similar ways in which these words can be used: one can be "chocolate-crazed" or "chocolate-addicted" and suffer "chocolate-induced" symptoms (e.g., each of these uses can be found in the pages of Wikipedia). In a context that gives rise to these expressions, it is unsurprising that chocolate should appear altogether more similar to a harmful narcotic.

In this paper we computationally model this idea that language use reflects category structure. As noted by De Leenheer and de Moor (2005), ontologies are lexical representations of concepts, so we can expect the effects of context on language use to closely reflect the effects of context on ontological structure. An understanding of the linguistic effects of context, as expressed through syntagmatic patterns of word usage, should lead therefore to the design of more flexible lexical ontologies that naturally adapt to their contexts of use. WordNet (Fell-

Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics, pages 57–64,
Prague, Czech Republic, June 2007.

baum, 1998) is just one such lexical ontology that can benefit greatly from the added flexibility that context-sensitivity can bring. Though comprehensive in scale and widely used, WordNet suffers from an obvious structural rigidity in which concepts are either entirely within a category or entirely outside a category: no gradation of category membership is allowed, and no contextual factors are brought to bear on criteria for membership. Thus, a gun is always a weapon in WordNet while an axe is never so, despite the uses (sporting or murderous) to which each can be put.

In section two we describe a computational framework for giving WordNet senses a functional, context-sensitive form. These functional forms simultaneously represent i) an intensional definition for each word sense; ii) a structured query capable of retrieving instances of the corresponding category from a context-specific corpus; and iii) a membership function that assigns gradated scores to these instances based on available syntagmatic evidence. In section three we describe how the knowledge required to automate this functional re-definition is acquired from the web and linked to WordNet. In section four we describe how these re-definitions can produce a robust model of metaphor, before we evaluate the descriptive sufficiency of this approach in section five, comparing it to the knowledge already available within WordNet. We conclude with some final remarks in section six.

2 Functional Context-Sensitive Categories

We take a wholly textual view of context and assume that a given context can be implicitly characterized by a representative text corpus. This corpus can be as large as a text archive or an encyclopedia (e.g., the complete text of Wikipedia), or as small as a single document, a sentence or even a single noun-phrase. For instance, the micro-context "alcoholic apple-juice" is enough to implicate the category Liquor, rather than Juice, as a semantic head, while "lovable snake" can be enough of a context to locally categorize Snake as a kind of Pet. There is a range of syntagmatic patterns that one can exploit to glean category insights from a text. For instance, the "X kills" pattern is enough to categorize X as a kind of Killer, "hunts X" is enough to categorize X as a kind of Prey, while "X-covered", "X-dipped" and "X-frosted" all indicate that X is a kind of Covering. Likewise, "army of X" suggests that a context views X as a kind of Soldier, while "barrage of X" suggests that X should be seen as a kind of Projectile.

We operationalize the collocation-type of adjective and noun via the function (*attr ADJ NOUN*), which returns a number in the range 0...1; this represents the extent to which ADJ is used to modify NOUN in the context-defining corpus. Dice's coefficient (e.g., see Cimiano *et al*., 2005) is used to implement this measure. A context-sensitive category membership function can be defined, as in that for Fundamentalist in Figure 1:

```
(define Fundamentalist.0 (arg0)
    (*  (max
            (%isa arg0 Person.0)
            (%isa arg0 Group.0))
        (min
            (max
                (attr political arg0)
                (attr religious arg0))
            (max
                (attr extreme arg0)
                (attr violent arg0)
                (attr radical arg0)))))
```

Figure 1. A functional re-definition of the category Fundamentalist.

The function of Figure 1 takes, as a single argument $\mathbf{arg}_0$, a putative member of the category Fundamentalist.0 (note how the sense tag, 0, is used to identify a specific WordNet sense of "fundamentalist"), and returns a membership score in the range 0...1 for this term. This score reflects the syntagmatic evidence for considering $\mathbf{arg}_0$ to be political or religious, as well as extreme or violent or radical. The function *(%isa* $\mathbf{arg}_0$ *CAT)* returns a value of 1.0 if some sense of $\mathbf{arg}_0$ is a descendent of CAT (here Person.0 or Group.0), otherwise 0. This safeguards ontological coherence and ensures that only kinds of people or groups can ever be considered as fundamentalists.

The example of Figure 1 is hand-crafted, but a functional form can be assigned automatically to many of the synsets in WordNet by heuristic means.

For instance, those of Figure 2 are automatically derived from WordNet's morpho-semantic links:

```
(define Fraternity.0 (arg_0)
    (*  (%sim arg_0 Fraternity.0)
        (max
            (attr fraternal arg_0)
            (attr brotherly arg_0))))

(define Orgasm.0 (arg_0)
    (*  (%sim arg_0 Orgasm.0)
        (max
            (attr climactic arg_0)
            (attr orgasmic arg_0))))
```

Figure 2. Exploiting the WordNet links between nouns and their adjectival forms.

The function *(%sim* $\boldsymbol{arg}_0$ *CAT)* reflects the perceived similarity between the putative member arg_0 and a synset CAT in WordNet, using one of the standard formulations described in Budanitsky and Hirst (2006). Thus, any kind of group (e.g., a glee club, a Masonic lodge, or a barbershop quartet) described in a text as "fraternal" or "brotherly" (both occupy the same WordNet synset) can be considered a Fraternity to the corresponding degree, tempered by its *a priori* similarity to a Fraternity; likewise, any climactic event can be categorized as an Orgasm to a more or less degree.

Alternately, the function of Figure 3 is automatically obtained for the lexical concept Espresso by shallow parsing its WordNet gloss: "strong black coffee brewed by forcing steam under pressure through powdered coffee beans".

```
(define Espresso.0 (arg_0)
    (*  (%sim arg_0 Espresso.0)
        (min
            (attr strong arg_0)
            (attr black arg_0))))
```

Figure 3. A functional re-definition of the category Espresso based on its WordNet gloss.

It follows that any substance (e.g., oil or tea) described locally as "black" and "strong" with a non-zero taxonomic similarity to coffee can be considered a kind of Espresso.

Combining the contents of WordNet 1.6 and WordNet 2.1, 27,732 different glosses (shared by 51,035 unique word senses) can be shallow parsed to yield a definition of the kind shown in Figure 3. Of these, 4525 glosses yield two or more properties that can be given functional form via *attr*. However, one can question whether these features are sufficient, and more importantly, whether they are truly diagnostic of the categories they are used to define. In the next section we consider another source of diagnostic properties, explicit similes on the web, before, in section 5, comparing the quality of these properties to those available from WordNet.

3 Diagnostic Properties on the Web

We employ the *Google* search engine as a retrieval mechanism for acquiring the diagnostic properties of categories from the web, since the *Google* API and its support for the wildcard term * allows this process to be fully automated. The guiding intuition here is that looking for explicit similes of the form "X is as P as Y" is the surest way of finding the most salient properties of a term Y; with other syntagmatic patterns, such as adjective:noun collocations, one cannot be sure that the adjective is central to the noun.

Since we expect that explicit similes will tend to exploit properties that occupy an exemplary point on a scale, we first extract a list of antonymous adjectives, such as "hot" or "cold", from WordNet. For every adjective ADJ on this list, we send the query "*as ADJ as **" to Google and scan the first 200 snippets returned to extract different noun values for the wildcard *. From each set of snippets we can also ascertain the relative frequencies of different noun values for ADJ. The complete set of nouns extracted in this way is then used to drive a second phase of the search, in which the query template "*as * as a NOUN*" is used to acquire similes that may have lain beyond the 200-snippet horizon of the original search, or that may hinge on adjectives not included on the original list. Together, both phases collect a wide-ranging series of core samples (of 200 hits each) from across the web, yielding a set of 74,704 simile instances (of 42,618 unique types) relating

3769 different adjectives to 9286 different nouns

3.1 Property Filtering

Unfortunately, many of these similes are not sufficiently well-formed to identify salient properties. In many cases, the noun value forms part of a larger noun phrase: it may be the modifier of a compound noun (as in "bread lover"), or the head of complex noun phrase (such as "gang of thieves" or "wound that refuses to heal"). In the former case, the compound is used if it corresponds to a compound term in WordNet and thus constitutes a single lexical unit; if not, or if the latter case, the simile is rejected. Other similes are simply too contextual or under-specified to function well in a null context, so if one must read the original document to make sense of the simile, it is rejected. More surprisingly, perhaps, a substantial number of the retrieved similes are ironic, in which the literal meaning of the simile is contrary to the meaning dictated by common sense. For instance, "as hairy as a bowling ball" (found once) is an ironic way of saying "as hairless as a bowling ball" (also found just once). Many ironies can only be recognized using world knowledge, such as "as sober as a Kennedy" and "as tanned as an Irishman".

Given the creativity involved in these constructions, one cannot imagine a reliable automatic filter to safely identify bona-fide similes. For this reason, the filtering task is performed by a human judge, who annotated 30,991 of these simile instances (for 12,259 unique adjective/noun pairings) as non-ironic and meaningful in a null context; these similes relate a set of 2635 adjectives to a set of 4061 different nouns. In addition, the judge also annotated 4685 simile instances (of 2798 types) as ironic; these similes relate a set of 936 adjectives to a set of 1417 nouns. Perhaps surprisingly, ironic pairings account for over 13% of all annotated simile instances and over 20% of all annotated types.

3.2 Linking to WordNet Senses

To create functional WordNet definitions from these adjective:noun pairings, we first need to identify the WordNet sense of each noun. For instance, "as stiff as a zombie" might refer either to a re-animated corpse or to an alcoholic cocktail (both are senses of "zombie" in WordNet, and drinks can be "stiff" too). Disambiguation is trivial for nouns with just a single sense in WordNet. For nouns with two or more fine-grained senses that are all taxonomically close, such as "gladiator" (two senses: a boxer and a combatant), we consider each sense to be a suitable target. In some cases, the WordNet gloss for as particular sense will literally mention the adjective of the simile, and so this sense is chosen. In all other cases, we employ a strategy of mutual disambiguation to relate the noun vehicle in each simile to a specific sense in WordNet. Two similes "as A as N_1" and "as A as N_2" are mutually disambiguating if N_1 and N_2 are synonyms in WordNet, or if some sense of N_1 is a hypernym or hyponym of some sense of N_2 in WordNet. For instance, the adjective "scary" is used to describe both the noun "rattler" and the noun "rattlesnake" in bona-fide (non-ironic) similes; since these nouns share a sense, we can assume that the intended sense of "rattler" is that of a dangerous snake rather than a child's toy. Similarly, the adjective "brittle" is used to describe both saltines and crackers, suggesting that it is the bread sense of "cracker" rather than the hacker, firework or hillbilly senses (all in WordNet) that is intended.

These heuristics allow us to automatically disambiguate 10,378 bona-fide simile types (85%), yielding a mapping of 2124 adjectives to 3778 different WordNet senses. Likewise, 77% (or 2164) of the simile types annotated as ironic are disambiguated automatically. A remarkable stability is observed in the alignment of noun vehicles to WordNet senses: 100% of the ironic vehicles always denote the same sense, no matter the adjective involved, while 96% of bona-fide vehicles always denote the same sense. This stability suggests two conclusions: the disambiguation process is consistent and accurate; but more intriguingly, only one coarse-grained sense of any word is likely to be sufficiently exemplary of some property to be useful in a simile.

4 From Similes to Category Functions

As noted in section 3, the filtered web data yields 12,259 bona-fide similes describing 4061 target nouns in terms of 2635 different adjectival properties. Word-sense disambiguation allows 3778 synsets in WordNet to be given a functional re-definition in terms of 2124 diagnostic properties, as in the definition of Gladiator in Figure 4:

*(define Gladiator.0 (**arg**$_0$)*
((%isa **arg**$_0$ Person.0)*
((%sim **arg**$_0$ Gladiator.0)*
(combine
*(attr strong **arg**$_0$)*
*(attr violent **arg**$_0$)*
*(attr manly **arg**$_0$)))))*

Figure 4. A web-based definition of Gladiator.

Since we cannot ascertain from the web data which properties are necessary and which are collectively sufficient, we use the function *combine* to aggregate the available evidence. This function implements a naïve probabilistic *or*, in which each piece of syntagmatic evidence is naively assumed to be independent, as follows:

(combine e_0 e_1) $= e_0 + e_1(1 - e_0)$
(combine e_0 $e_1 \ldots e_n$) = (combine e_0 (combine $e_1 \ldots e_n$))

Thus, any combatant or competitor (such as a sportsman) that is described as *strong*, *violent* or *manly* in a corpus can be categorized as a Gladiator in that context; the more properties that hold, and the greater the degree to which they hold, the greater the membership score that is assigned.

The source of the hard taxonomic constraint *(%isa* arg_0 *Person.0)* is explained in the next section. For now, note how the use of *%sim* in the functions of Figures 2, 3 and 4 means that these membership functions readily admit both literal and metaphoric members. Since the line between literal and metaphoric uses of a category is often impossible to draw, the best one can do is to accept metaphor as a gradable phenomenon (see Hanks, 2006). The incorporation of taxonomic similarity via *%sim* ensures that literal members will tend to receive higher membership scores, and that the most tenuous metaphors will receive the lowest membership scores (close to 0.0).

4.1 Constrained Category Inclusion

Simile and metaphor involve quite different conceptual mechanisms. For instance, anything that is particularly strong or black might meaningfully be called ”as black as espresso” or ”as strong as espresso”, yet few such things can meaningfully be called just ”espresso”. While simile is a mechanism for highlighting inter-concept similarity, metaphor is at heart a mechanism of category inclusion (see Glucksberg, 2001). As the espresso example demonstrates, category inclusion is more than a matter of shared properties: humans have strong intuitions about the structure of categories and the extent to which they can be stretched to include new members. So while it is sensible to apply the category Espresso to other substances, preferably liquids, it seems nonsensical to apply the category to animals, artifacts, places and so on.

Much as the salient properties of categories can be acquired form the web (see section 3), so too can the intuitions governing inclusion amongst categories. For instance, an attested web-usage of the phrase ”Espresso-like CAT” tells us that sub-types of CAT are allowable targets of categorization by the category Espresso. Thus, since the query ”espresso-like substance” returns 3 hits via Google, types of substance (oil, etc.) can be described as Espresso if they are contextually strong and black. In contrast, the query ”espresso-like person” returns 0 hits, so no instance of person can be described as Espresso, no matter how black or how strong. While this is clearly a heuristic approach to a complex cognitive problem, it does allow us to tap into the tacit knowledge that humans employ in categorization. More generally, a concept X can be included in a category C if X exhibits salient properties of C and, for some hypernym H of X in WordNet, we can find an attested use of ”C-like H” on the web.

If we can pre-fetch all possible ”C-like H” from the web, this will allow comprehension to proceed without having to resort to web analysis in mid-categorization. While there are too many possible values of H to make full pre-fetching a practical reality, we *can* generalize the problem somewhat, by selecting a range of values for H from the middle-layer of WordNet, such as *Person, Substance, Animal, Tool, Plant, Structure, Event, Vehicle, Idea* and *Place*, and by pre-fetching the query ”C-like H” for all 4061 nouns collected in section 3, combined with this limited set of H values. For every noun in our database then, we pre-compile a vector of possible category inclusions.

For instance, "lattice" yields the following vector:

{*structure(1620), substance(8), container(1), vehicle(1)*}

where numbers in parentheses indicate the web-frequency of the corresponding "Lattice-like H" query. Thus, the category Lattice can be used to describe (and metaphorically include) other kinds of structure (like crystals), types of substance (e.g., crystalline substances), containers (like honey-combs) and even vehicles (e.g., those with many compartments). Likewise, the noun "snake" yields the following vector of possibilities:

{*structure(125), animal(122), person(56), vehicle(17), tool(9)*}

(note, the frequency for "person" includes the frequency for "man" and "woman"). The category Snake can also be used to describe and include structures (like tunnels), other animals (like eels), people (e.g., the dishonest variety), vehicles (e.g., articulated trucks, trains) and tools (e.g., hoses). The noun "gladiator" yields a vector of just one element, {*person(1)*}, from which the simple constraint *(%isa* arg_0 *Person.0)* in Figure 4 is derived. In contrast, "snake" is now given the definition of Figure 5:

```
(define Snake.0 (arg0)
    (*  (max
            (%isa arg0 Structure.0)
            (%isa arg0 Animal.0)
            (%isa arg0 Person.0)
            (%isa arg0 Vehicle.0))
        (* (%sim arg0 Snake.0)
          (combine
                (attr cunning arg0)
                (attr slippery arg0)
                (attr flexible arg0)
                (attr slim arg0)
                (attr sinuous arg0)
                (attr crooked arg0)
                (attr deadly arg0)
                (attr poised arg0)))))
```

Figure 5. A membership function for Snake using web-derived category-inclusion constraints.

Glucksberg (2001) notes that the same category, used figuratively, can exhibit different qualities in different metaphors. For instance, Snake might describe a kind of crooked person in one metaphor, a poised killer in another metaphor, and a kind of flexible tool in yet another. The use of *combine* in Figure 5 means that a single category definition can give rise to each of these perspectives in the appropriate contexts. We therefore do not need a different category definition for each metaphoric use of Snake.

To illustrate the high-level workings of category-inclusion, Table 1 generalizes over the set of 3778 disambiguated nouns from section 3 to estimate the propensity for one semantic category, like Person, to include members of another category, like Animal, in X-like Y constructs.

X-like Y	*P*	*A*	*Sub*	*T*	*Str*
(P)erson	.66	.05	.03	.04	.09
(A)nimal	.36	.27	.04	.05	.15
(Sub)stance	.14	.03	.37	.05	.32
(T)ool	.08	.03	.07	.22	.34
(Str)ucture	.04	.03	.03	.03	.43

Table 1. The Likelihood of a category X accommodating a category Y.

Table 1 reveals that 36% of "ANIMAL-like" patterns on the web describe a kind of Person, while only 5% of "PERSON-like" patterns on the web describe a kind of Animal. Category inclusion appears here to be a conservative mechanism, with like describing like in most cases; thus, types of Person are most often used to describe other kinds of Person (comprising 66% of "PERSON-like" patterns), types of substance to describe other substances, and so on. The clear exception is Animal, with "ANIMAL-like" phrases more often used to describe people (36%) than other kinds of animal (27%). The anthropomorphic uses of this category demonstrate the importance of folk-knowledge in figurative categorization, of the kind one is more likely to find in real text, and on the web (as in section 3), rather than in resources like WordNet.

5 Empirical Evaluation

The simile gathering process of section 3, abetted by Google's practice of ranking pages according to popularity, should reveal the most frequently-used comparative nouns, and thus, the most useful categories to capture in a general-purpose ontology like WordNet. But the descriptive sufficiency of these categories is not guaranteed unless the defining properties ascribed to each can be shown to be collectively rich enough, and individually salient enough, to predict how each category is perceived and applied by a language user.

If similes are indeed a good basis for mining the most salient and diagnostic properties of categories, we should expect the set of properties for each category to accurately predict how the category is perceived as a whole. For instance, humans – unlike computers – do not generally adopt a dispassionate view of ideas, but rather tend to associate certain positive or negative feelings, or affective values, with particular ideas. Unsavoury activities, people and substances generally possess a negative affect, while pleasant activities and people possess a positive affect. Whissell (1989) reduces the notion of affect to a single numeric dimension, to produce a *dictionary of affect* that associates a numeric value in the range 1.0 (most unpleasant) to 3.0 (most pleasant) with over 8000 words in a range of syntactic categories (including adjectives, verbs and nouns). So to the extent that the adjectival properties yielded by processing similes paint an accurate picture of each category / noun-sense, we should be able to predict the affective rating of each vehicle via a weighted average of the affective ratings of the adjectival properties ascribed to these nouns (i.e., where the affect rating of each adjective contributes to the estimated rating of a noun in proportion to its frequency of co-occurrence with that noun in our simile data). More specifically, we should expect that ratings estimated via these simile-derived properties should correlate well with the independent ratings contained in Whissell's dictionary.

To determine whether similes do offer the clearest perspective on a category's most salient properties, we calculate and compare this correlation using the following data sets:

A. Adjectives derived from annotated bona-fide (non-ironic) similes only.

B. Adjectives derived from all annotated similes (both ironic and non-ironic).

C. Adjectives derived from ironic similes only.

D. All adjectives used to modify a given noun in a large corpus. We use over 2-gigabytes of text from the online encyclopaedia Wikipedia as our corpus.

E. The set of 63,935 unique property-of-noun pairings extracted via shallow-parsing from WordNet glosses in section 2, e.g., *strong* and *black* for Espresso.

Predictions of affective rating were made from each of these data sources and then correlated with the ratings reported in Whissell's dictionary of affect using a two-tailed Pearson test ($p < 0.01$). As expected, property sets derived from bona-fide similes only (A) yielded the best correlation (+0.514) while properties derived from ironic similes only (C) yielded the worst (-0.243); a middling correlation coefficient of 0.347 was found for all similes together, demonstrating the fact that bona-fide similes outnumber ironic similes by a ratio of 4 to 1. A weaker correlation of 0.15 was found using the corpus-derived adjectival modifiers for each noun (D); while this data provides quite large property sets for each noun, these properties merely reflect potential rather than intrinsic properties of each noun and so do not reveal what is most diagnostic about a category. More surprisingly, property sets derived from WordNet glosses (E) are also poorly predictive, yielding a correlation with Whissell's affect ratings of just 0.278. This suggests that the properties used to define categories in hand-crafted resources like WordNet are not always those that actually reflect how humans think of these categories.

6 Concluding Remarks

Much of what we understand about different categories is based on tacit and defeasible knowledge of the outside world, knowledge that cannot easily be shoe-horned into the rigid *is-a* structure of an ontology like WordNet. This already-complex picture is complicated even further by the often metaphoric relationship between words and the categories they denote, and by the fact that the metaphor/literal distinction is not binary but gradable. Furthermore, the gradability of category membership is clearly influenced by context: in a corpus describing the exploits of Vikings, an axe will most likely be seen as a kind of weapon, but in a corpus dedicated to forestry, it will likely describe a tool. A resource like WordNet, in which *is-a* links are reserved for category relationships that are always true, in any context, is going to be inherently limited when dealing with real text.

We have described an approach that can be seen as a functional equivalent to the CPA (Corpus Pattern Analysis) approach of Pustejovsky *et al*. (2004), in which our goal is not that of automated induction of word senses in context (as it is in CPA) but the automated induction of flexible, context-sensitive category structures. As such, our goal is primarily ontological rather than lexicographic, though both approaches are complementary since each views syntagmatic evidence as the key to understanding the use of lexical concepts in context. By defining category membership in terms of syntagmatic expectations, we establish a functional and gradable basis for determining whether one lexical concept (or synset) in WordNet deserves to be seen as a descendant of another in a particular corpus and context. Augmented with ontological constraints derived from the usage of ”X-like Y” patterns on the web, we also show how these membership functions can implement Glucksberg’s (2001) theory of category inclusion.

We have focused on just one syntagmatic pattern here – adjectival modification of nouns – but categorization can be inferred from a wide range of productive patterns in text, particularly those concerning verbs and their case-fillers. For instance, verb-centred similes of the form ”to V+*inf* like a|an N” and ”to be V+*past* like a|an N” reveal insights into the diagnostic behaviour of entities (e.g., that predators hunt, that prey is hunted, that eagles soar and bombs explode). Taken together, adjective-based properties and verb-based behaviours can paint an even more comprehensive picture of each lexical concept, so that e.g., political agents that kill can be categorized as assassins, loyal entities that fight can be categorized as soldiers, and so on. An important next step, then, is to mine these behaviours from the web and incorporate the corresponding syntagmatic expectations into our category definitions. The symbolic nature of the resulting definitions means these can serve not just as mathematical membership functions, but as ”active glosses”, capable of recruiting their own members in a particular context while demonstrating a flexibility with categorization and a genuine competence with metaphor.

References

Alexander Budanitsky and Graeme Hirst. 2006. Evaluating WordNet-based Measures of Lexical Semantic Relatedness. *Computational Linguistics*, 32(1), pp 13-47.

Christiane Fellbaum (ed.). 1998. *WordNet: An Electronic Lexical Database*. The MIT Press, Cambridge, MA.

Cynthia Whissell. 1989. The dictionary of affect in language. In R. Plutchnik & H. Kellerman (Eds.). *Emotion: Theory and research*. New York, Harcourt Brace, 113-131.

James Pustejovsky, Patrick Hanks and Anna Rumshisky. 2004. Automated Induction of Sense in Context. In Proceedings of COLING 2004, Geneva, pp 924-931.

Patrick Hanks. 2006. Metaphoricity is a Gradable. In A. Stefanowitsch and S. Gries (eds.). *Corpora in Cognitive Linguistics. Vol. 1: Metaphor and Metonymy*. Berlin: Mouton.

Patrick Hanks. 2004. The syntagmatics of metaphor and idiom. *International Journal of Lexicography*, 17(3).

Philipp Cimiano, Andreas Hotho, and Steffen Staab. 2005. Learning Concept Hierarchies from Text Corpora using Formal Concept Analysis. *Journal of AI Research*, 24: 305-339.

Pieter De Leenheer and Aldo de Moor. 2005. Context-driven Disambiguation in Ontology Elicitation. In Shvaiko P. & Euzenat J. (eds.), Context and Ontologies: Theory, Practice and Applications, AAAI Tech Report WS-05-01. AAAI Press, pp 17-24.

Sam Glucksberg. 2001. *Understanding figurative language: From metaphors to idioms*. Oxford: Oxford University Press.