A Dynamic-Type Hierarchy for Linguistic Creativity

Tony Veale

Department of Computer Science,
University College Dublin,
Belfield, Dublin 6,
Ireland.

Tony.Veale@UCD.ie

Abstract

Metaphor and analogy are perhaps the most challenging aspects of linguistic creativity for a conceptual representation to facilitate, since by their very nature they seek to stretch the boundaries of domain description and dynamically establish new ways of determining inter-domain similarity. By solving the vexing representational problems posed by these phenomena, we can create a more fluid conceptual organization that is more suited to creative processing in general. Toward this end, this paper considers the problem of how a conceptual system structured around a central taxonomy can dynamically create new categories or types to understand creative metaphors and analogies. This theoretical perspective yields a practical method for dynamic type creation within the WordNet lexical knowledge-base that has applications to both creative writing systems and intelligent information retrieval.

Keywords: analogy, type hierarchies, WordNet, linguistic creativity, metaphor

1 Introduction

Linguistic creativity may sometimes seem like superficial word-play, yet in its most potent guises it has the power to change the way we see and represent the world. In this paper we consider two of the most representationally-challenging linguistic armaments in our creative arsenal, metaphor and analogy. These processes are interesting to linguists and creativity researchers alike, because they often exploit latent similarities between domains that expose the holes in our mental lexicon and the structural inadequacies of our underlying conceptual system [13]. In particular, because metaphors and analogies are used to create new ways of thinking about familiar things, they reveal the fluid boundaries that exist both between the conceptual categories we use to structure the world [8] and the words we use to communicate these categories [14]. This fluidity contrasts sharply with the rigidity of the taxonomies that have been traditionally posited to organize our category systems [7, 11, 22]. However, there is no need to replace these rigid taxonomies outright; our goal in this paper is to show how fluidity can be considered an emergent property of existing taxonomies when taxonomic types are treated as dynamic, rather than static entities that are created as they are needed.

Taxonomies have, since antiquity [7], provided a systematic means of hierarchical decomposition of knowledge, allowing a domain to be successively dissected via differentiation into smaller pockets of related concepts. Taxonomic differentiation leads to effective clustering, so that similar concepts are situated in the same region of the taxonomy. This locality of meaning not only makes the similarity of different categories easier to assess computationally, it also means that the elements of a domain tend to be clustered around the same parent types, which can thus act as indices into the domain for effective analogical mapping. Indeed, the first account of metaphor as a conceptual process, as offered by Aristotle in his *Poetics* [7], was wholly taxonomic. In the Aristotelian scheme, two concepts can be metaphorically or analogically connected if a common taxonomic parent can be found to unify them both. The crucial role of a central taxonomic backbone in organizing knowledge survives today in such large-scale ontologies as Cyc [9], a common-sense ontology for general reasoning, and WordNet [11], a psycholinguistically-motivated and very comprehensive database of lexical concepts in English. The Aristotelian view of taxonomic metaphor also continues to exert considerable influence in computational theories, as shown by [5] and [22].

Yet, if a taxonomy is to be a driving force in the understanding of metaphor and analogy, it must anticipate every possible point of comparison between every pair of domains. However, to even suggest that such an exhaustive taxonomy is possible – and the idea certainly raises grave concerns about tractability – would be to diminish the role of metaphor and analogy as tools for affecting change to our conceptual systems. To resolve this contradiction, authors such as Eileen Cornell Way [22] have argued for the importance of a dynamic type hierarchy (DTH) as a taxonomic backbone for conceptual structure. Such a taxonomy would dynamically reorganize itself to reveal new types in response to appropriate metaphors. For example, Way [22] gives as an example "Nixon is the submarine of world politics", and suggests that this metaphor is resolved by the dynamic type *ThingsWhichBehaveInASecretOrHiddenManner*. However, as useful as a dynamic hierarchy would be for metaphor and analogy, Way does not suggest an empirical means of constructing a DTH capable of generating such ambitious types, leaving the issue of exhaustiveness, and all it entails for computational tractability, unresolved.

This paper describes an automated means of constructing a DTH that dynamically generates new taxonomic types in response to creative analogies and metaphors. The underlying static type hierarchy is provided by the taxonomic structure of WordNet [11], while dynamic types are extracted when needed from the flat textual glosses that annotate

every concept in WordNet (see [6] for comparable uses of these glosses as sources of tacit structure). To achieve this dynamism, we identify a taxonomic meta-type we call an "analogical pivot" that facilitate the processes of analogue retrieval and mapping [15] in a taxonomy, and show how types in existing taxonomies like WordNet and Cyc, which contain relatively few such pivots naturally, are automatically converted into pivots as new sub-types are dynamically added to the taxonomy. We intend to demonstrate that that the conventional wisdom regarding metaphor and analogy – that such processes are creative because clever word-play is indicative of an underlying mental agility and suppleness of conceptual structure – also withstands theoretical scrutiny when considered from the perspective of current creativity research.

2 Analogical Pivots

Taxonomic systematicity requires that analogous domains should be differentiated in the same ways, so that similarity judgments in each domain will be comparable. But in very large taxonomies, this systematicity is often lacking. Consider WordNet, whose taxonomy is designed to mirror the lexical structure of the English language. WordNet is a directed-graph of semantically-related synonym sets or "synsets", where each synset is a collection of synonymous words that collectively identify a single concept. For instance, the synset *{letter, alphabetic_character}* denotes the alphabetic rather than the epistolary sense of "letter". WordNet reflects our linguistic intuitions of how concepts should be structured, but the sheer scale of the task ensures that this is not always achieved in a complete and consistent manner. Thus, while the concept *{alphabet}* is differentiated along cultural lines into the hyponyms *{Greek_alphabet}* and *{Hebrew_alphabet}*, the concept *{letter, alphabetic_character}* is not similarly differentiated into *{Greek_letter}* and *{Hebrew_letter},* even though there are specific letter concepts that would benefit from this organization. Instead, every letter of each alphabet, such as *{alpha}* and *{aleph}*, is located under exactly the same hypernym, *{letter, alphabetic_character}*. This means that on structural grounds alone, each letter is equally similar to every other letter, no matter what alphabet they belong to (e.g., *{alpha}* is as similar to *{aleph}* as it is to *{beta}*). Crucially, more than similarity judgment is impaired, for a lack of systematicity and symmetry in differentiation undermines another core rationale of taxonomic structure, the ability to recognize analogies and metaphors. Thus, the structure of WordNet is not sufficiently differentiated to facilitate either the generation or the comprehension of analogies such as "what is the Jewish gamma?" *(gimel)* or "who is the Viking Ares" *(Tyr)*. This kind of analogical task provides an acid-test for the ability to a representation to support creative reasoning.

To see what needs to be added to WordNet, consider the analogical compound "Hindu Zeus" and how one might interpret it using the existing structure of WordNet. The goal is to find a counterpart for the source concept Zeus (the supreme deity of the Greek pantheon) in the target domain of Hinduism. In WordNet 1.6, *{Zeus}* is a daughter of *{Greek_deity}*, which is turn is a daughter of *{deity, god}*. A fragment of the WordNet taxonomy that represents deities is illustrated in Figure 1; hyponymy links are represented as arrows, and for convenience, each link is annotated with the salient property of the hyponym. These annotations are not explicitly represented in WordNet, but can be found as part of the textual gloss associated with each synset. As we shall discuss later, the extraction of these highly specific properties forms a key part of the DTH construction process.

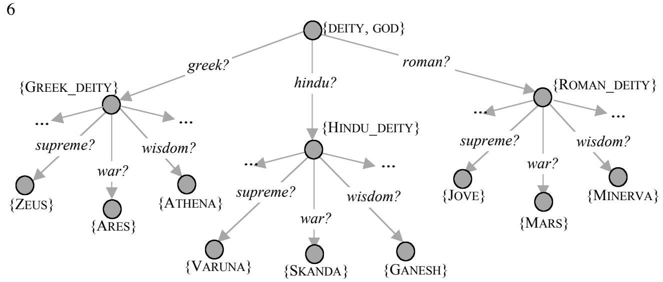

Figure 1. A decision-tree perspective on the sub-taxonomy of *{deity, god}* in WordNet.

If WordNet were explicitly annotated or signposted in this way, then the analogical denotation of "Hindu Zeus" would be straightforward to find. The signposted taxonomy in Figure 1 reveals that *{Zeus}* is a "supreme" variant of a *{Greek_deity}*, and that *{Greek_deity}* is a "Greek" variant of *{deity, god}*. Since *{Hindu_deity}* is also a hyponym of *{deity, god}*, this represents the first opportunity to move from the domain of "Greek" deities to the domain of "Hindu" counterparts, by essentially following the "Hindu" signpost from *{deity, god}* to *{Hindu_deity}*. We thus identify the type *{deity, god}* as the analogical pivot around which the mapping turns. Drilling further into the Hindu domain, we see that only *{Varuna}* is signposted as being a "supreme" variant of *{Hindu_deity}*. The concept *{Varuna}* is therefore an ideal, unambiguous mapping for the source concept *{Zeus}*.

Signposts are shown in Figure 1 as explicit labels on hyponym/hypernym links, but notice how WordNet already internalizes these labels as modifier terms in compound types like *{Greek_deity}*, *{Hindu_deity}* and *{Roman_deity}*. We find then in WordNet a taxonomy that is already partially signposted, where compound types like these are used as a middle layer of indexing structure to mediate between very broad stratum of superordinate concepts (like *{deity, god}*) and a very specific stratum of instances (like *{Zeus}*). To make the deities taxonomy of Figure 1 completely self-signposting in this way, we simply need to add the compound types *{supreme_deity}, {wisdom_deity}* and *{war_deity}* to this middle layer. If these additions can be automated based on the existing structure of the taxonomy and the contents of the available glosses, then in principle the entire WordNet taxonomy can become self-signposting. We consider an algorithmic means to achieve this goal in section 3.

Given a sufficiently signposted taxonomy, the following serves as a general method for resolving metaphors and analogies: given a source concept and a term denoting a target domain, climb up the taxonomy from the source concept until a hypernym, called the pivot, is encountered. The pivot is any superordinate type from which a signpost leads to the desired target domain. In general, if the pivot has the form {P} and the target domain is alluded to by the term "T", then "T" serves as a signpost from {P} to {T_P}. Computationally, the mechanism is simply one of lexical concatenation: once {P} is identified, then T + P = "T_P" denotes the target category {T_P} in which to search for a more precise target mapping. Again, the pivot {P} is computationally simple to identify – it is any hyponym of the source concept for which T + P = "T_P" denotes an existing WordNet concept {T_P}.

The basic approach as outlined above is highly sensitive to surface variations in the expression of the analogy. It works well for "Hindu Zeus" because "Hindu" + *{deity, god}* denotes an existing WordNet concept *{Hindu_deity}*. However, were the analogy to be expressed as "Hindoo Zeus" or "Hindustani Zeus", or even "Trimurti Zeus" or "Brahman Zeus", no pivot would be identified for the analogy as concatenation would produce, in each case, a target category that simply does not exist in WordNet (i.e., the WordNet synset *{Hindu_deity}* is a singleton and does not contain *Hindoo_deity* or any other variation). We thus need to complicate the approach a little, and introduce the possibility that the target concept will be identified instead by T' + P = "T'_P" → {T'_P}, where T' is any lexical variant or proxy of "T" such as a synonym or a metonym.

2.1 Classical Approaches

This use of generalization and specialization is clearly reminiscent of the classic Aristotelian approach. Aristotle's theory of metaphor and proportional analogy, as expressed in the *Poetics* [7], suggests that two entities or taxonomic types can be considered analogous if they share a common superordinate type. In fact, Aristotle's approach to analogy still finds considerable traction today, in computational models that seek to unify concepts through generalization (e.g., see [5, 22]), but its simplicity means that it is easily trivialized. In a well designed taxonomic hierarchy, any two concepts will always share at least one superordinate type, even if it is the root type, and so any two concepts will always be potential analogues in such a system. The signposting metaphor allows us to see why the basic Aristotelian approach is deficient, since the approach can be said to use just two signposts: one labeled "UP", which allows indiscriminate generalization of the source concept, and one labeled "DOWN", which subsequently allows indiscriminate specialization into the target domain. The classical approach thus lacks resolving power, and is unable to differentiate among different target candidates in a mapping. The problem is that Aristotle's approach, and the computational models based on it, do not make sufficient use of the semantic implications of hyponymy and hypernymy to constrain how the taxonomy is traversed. The key to making Aristotle's basic intuitions about analogy computationally workable is the recognition that all hyponym links are not alike. Rather, each should be seen as a nuanced semantic relationship between types, so that rather than X "isa" Y we have e.g., X is the "supreme" version of Y.

3 Dynamic Types and Taxonomic Signposting

WordNet has not been designed as a case-base for analogical reasoning, but as a general purpose lexical resource. As a result of this generality, well-signposted pivots like *{deity,* *god}* are extremely rare in WordNet. For example, the concept *{letter, alphabetic_character}* should, in principle, be the ideal pivot. Reflecting the fact that *{alphabet}* is differentiated into *{Greek_alphabet}* and *{Hebrew_alphabet}*, it should be perched above a variety of domain variations like *{Greek_letter}* and *{Hebrew_letter},* where these in turn organize families of specific letter concepts, like *{alpha}, {beta}*, *{gamma}* and *{aleph}, {beth}, {gimel}*, into culturally signposted groupings. Unfortunately, WordNet does not contain this middle layer of domain variations, and so lacks the signposts to resolve analogies like "Jewish delta" or "Greek gimel". Of course, even if WordNet did contain these signpost concepts, which would allow *{letter, alphabetic_character}* to act as an analogical pivot, these specific signposts would only direct the mapping to the general vicinity of the best target concept and no further. WordNet would still lack the structural finesse to pick out the specific member of the target category suggested by the pivot, e.g., to determine that *{daleth}* is the most appropriate hyponym of *{Hebrew_letter}* to align with *{delta}* in the analogy "Jewish delta".

Figure 2 illustrates the representation of the *{letter, alphabetic_character}* domain as we find it in WordNet 1.6. Note how there is insufficient differentiating structure to discriminate letters in one alphabet from the those of another, or indeed, from other letters of the same alphabet, which makes a precise mapping of domain counterparts impossible.

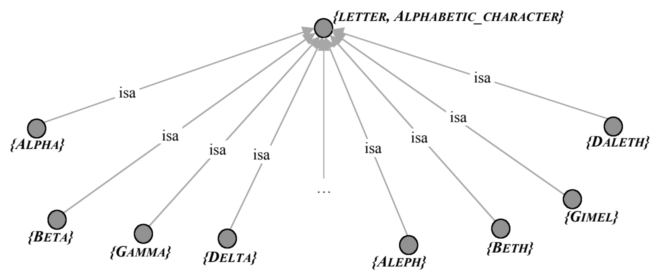

Figure 2. The impoverished *{letter, alphabetic_character}* taxonomy in WordNet

What is needed is two successive layers of signposts, one layer to connect the pivot to broad groupings of sub-types that share a general property (like "Greek"), and another, finer layer, to connect the pivot to even smaller groupings that each share a more specific property (like "occupies the fourth place in the alphabet"). The broad layer would allow the mapping process to identify the vicinity of the taxonomy in which the analogical target resides, while the finer layer would then allow the mapping to identify a specific concept in this vicinity. Such an arrangement of overlapping signposts is illustrated in Figure 3.

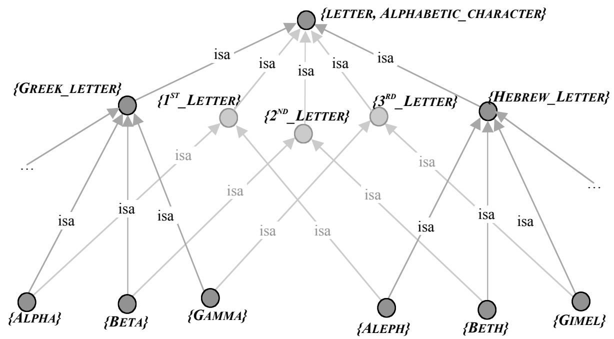

Figure 3. The taxonomic structure of *{letter, alphabetic_character}* becomes a rich lattice

In the idealized organization of Figure 3, the *{letter, alphabetic_character}* type is not only broadly differentiated into the cultural sub-types *{Greek_letter}* and *{Hebrew_letter}*, it is also finely differentiated by relative letter position into *{1st_letter}*, *{2nd_letter}* and so on. This structure is sufficient to allow a near-isomorphic mapping to be generated from one alphabet to another (with the exception of extra letters in one alphabet that have no true analogue in the other domain), by first mapping from one alphabet system to another, and then mapping from one relative position to another.

The goal of dynamic type creation is to turn the impoverished structure of Figure 2 into the signposted structure of Figure 3 as the need for these signposts is recognized during analogical reasoning. The presence of enough signposts allows nodes like *{deity, god}* and *{letter, alphabetic_character}* to be identified as pivots during analogical processing. When enough potential pivots are in place, a taxonomy like WordNet becomes a powerful decision-lattice for analogical reasoning. As it stands, the existing structure of WordNet can already be thought of as a partial implementation of a very general decision-lattice for all possible analogies. As illustrated in Figure 1, WordNet already contains some signposting structure, in the form of non-leaf compounds like *{Greek_deity}* and *{Hindu_deity}*, though it lacks far more than it contains. Fortunately, WordNet also contains a rich source of non-structural information in the collected textual glosses that accompany each synset. These glosses are handcrafted by lexicographers to be pithy and to the point, and like dictionary definitions, contain lexical terms that capture the important semantic dimensions of the associated concept. For instance, the WordNet 1.6 gloss for *{Zeus}* notes that Zeus is the "supreme" being in the Greek pantheon of gods. A dynamic type hierarchy can exploit the contents of these glosses to create new types, like *{supreme_deity}*, as they are needed. As more analogies are processed over time, a dynamic WordNet will go from the state illustrated in Figure 2 to that illustrated in Figure 3.

3.1 Type Creation As Feature Reification

WordNet, not unlike other taxonomies such as Cyc [9], expresses only some of its intended structure explicitly, via hyponymy links, and most of it implicitly, via handcrafted textual glosses. These glosses are primarily intended for human rather than machine consumption, but most are nonetheless consistent enough for the latter, enabling the automatic extraction of additional semantic features or relationships between concepts. WordNet glosses have effectively been mined in the past to automatically extract new lateral connections between concepts, that allow text comprehension and reasoning systems to perform

complex inferences (see [6] and [18]). To opportunistically build a dynamic type system on WordNet glosses, we need a means of recognizing those gloss terms that have the potential to yield the most discriminating analogical signposts. Once recognized, these words can be lifted out of the gloss and reified, through combination with a potential pivot type, to create new taxonomic types. For example, if the word "supreme" is recognized as having analogical potential in the gloss of *{Zeus}*, it can be lifted out and reified, via combination with the hyponym *{deity, god}*, to create a new type *{supreme_deity}*. The same process will ensue for *{Varuna}*, whose gloss also contains the term "supreme", so together, the signposts *{Hindu_deity}*, which already exists, and *{supreme_deity}*, which is dynamically created, can facilitate an unambiguous mapping from *{Zeus}* to *{Varuna}*. The reification process effectively anticipates the hypernym that will later serve as the pivot during an analogy. Since the immediate hypernym of *{Zeus}*, *{Greek_deity}* is already signposted as a "Greek" variant of *{deity, god}*, we do not choose this as the basis for reification. The most specific non-signposted hypernym of *{Zeus}* is *{deity, god}* and thus we use this to form the dynamic type: "supreme" + *{deity, god} = {supreme_deity}*.

The noun sense glosses of WordNet 1.6 collectively contain over 40,000 unique content words. It would not thus make sense to reify all the content words in the gloss of a concept, as this would result in an overcrowded taxonomy with many bizarre, meaningless or vague types. But neither would it make sense to handcraft a vocabulary of reifiable terms, especially since the analogical utility of reifying a term crucially depends on the location of the concept it annotates in the taxonomy. It happens that "supreme" is a good choice in the context of deities, as there exists a supreme deity in a number of pantheons, but "supreme" might not be so useful a discriminator in another domain, such as that of cars.

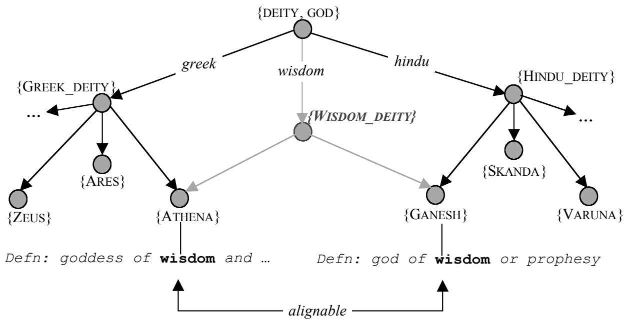

Figure 4. The glosses of *{Athena}* and *{Ganesh}* contain the reifiable feature *wisdom*.

To avoid reifying terms that are semantically vacuous and which have no true discriminating power, only those terms that have proven discrimination power should be reified. The pragmatics of compound noun formation ensure that if a term is already used as a modifier in an existing compound concept, then it non-trivially differentiates the head of that compound. Consider the word "wisdom", which occurs in 20 different WordNet glosses, which is enough to demonstrate it has a cross-over appeal between domains, but not too many to suggest that it is semantically indistinct. Additionally, there is at least one WordNet precedent, *{wisdom_tooth}*, for its use as a meaningful differentiator in another domain. This leads us to define differentiation potential as follows:

(1)

Definition: A lemmatized word-form has *differentiation potential* if it occurs in more than one gloss, but not in so many glosses (e.g., more than 1000) as to be semantically vacuous. There must also be a precedent for using this word-form as a modifier in at least one existing compound type

Looking to Figure 4, we see that "wisdom" occurs in the gloss of at least three concepts – *{Athena}*, *{Ganesh}* and *{Minerva}* – that are all descendent hyponyms of the type *{deity, god}*. Since each of these hyponyms resides at the same relative depth from this common hypernym, one can assume that each possesses the same degree of specificity. This suggests that it might be analogically useful to create a new taxonomic type *{wisdom_deity}* to unite these instances into a single category of their own. We can summarize this intuition as follows:

(2)

Definition: A word-form has *alignment potential* relative to a given hypernym type if it can be found in the glosses of multiple concepts that reside in different places in the taxonomy, at the same relative depth from a common hypernym

For a term to be worth reifying, it must possess both differentiation potential and alignment potential. So in principle, a term with differentiation potential (according to definition 1) in the gloss of a given concept can be used to signpost any number of hypernyms for which it has analogical potential (according to definition 2). The former constraint helps ensure that nonsense types are not created, while the latter ensures that a dynamic type is only created when there is at least one analogical pairing that can potentially be facilitated by it. It follows that the further up the taxonomy one adds new signposts via these new types, the greater the semantic distance that can be bridged by any future analogy.

4 Ontological Creativity

The conventional wisdom regarding the creative nature of metaphor and analogy is unequivocally clear: together they represent the acme of linguistic inventiveness and verbal intelligence. But does this intuition withstand theoretical scrutiny? In particular, how creative is the taxonomic approach to metaphor and analogy described in this paper when judged relative to the computational frameworks of current creativity research?

Creativity is conventionally modeled as an exploratory process in the space of possible concepts, to discover pockets of hitherto unrealized or unevaluated elements that exhibit both novelty and value [1,2,23]. If these elements are ascribed high-value but have properties that are characteristic of concepts with little or no value, then the discovery may be said to be both useful and surprising. The space may be defined explicitly, by some foundational axioms (which may themselves by modified to explore the space of possible spaces, in what is termed transformational creativity), or implicitly, by a set of existing conceptual structures from which other, unrealized structures can be generalized (see [3, 4, 17, 18, 19]). The process of dynamic type creation operates in the space of possible concepts of which a given ontology, such as WordNet, represents just a small lexically-realized sampling [11, 17, 18]. The extent of the space outside this sampling that can be explored when constructing a DTH is constrained both by this set of realized ontological entries, which are needed to serve as hypernyms for any new type, and the set of content words that occur in the glosses of these entries.

Though the effect of dynamic type creation is to transform the ontology, it would be inaccurate to describe this dynamism as transformational creativity. The underlying space of possible ontological entries is not itself transformed by type creation, it is simply the case that elements in the space are moved from a state of unrealized possibility to a state of realized actuality. Nonetheless, if the ontology is not viewed as the creative space itself, as it is here, but as the vocabulary through which an even larger space is defined, dynamic type creation at the ontological level can meaningfully be considered transformational. For example, if the ontology is used to define the space of all composite conceptual structures, as it might in models of conceptual blending [15, 16], then additions to the ontology would have transformational effects on the structure of this compositional space.

It may be useful to consider the dynamic types in a DTH as the result of a creative induction process [4]. In general, creative induction allows new concepts to be induced to facilitate the solution of a problem that could not be achieved within a given logical framework without those new concepts. The DTH described here conforms to this general pattern of need-driven concept creation, even if the concepts so created are for the most part P- rather than H-creative [1,2]. For example, given the ontological problem of measuring the deep similarity between two taxonomically distant concepts, a DTH can induce new types on the basis of specific ontological evidence (a combination of gloss terms and hypernyms) to reveal that these concepts are as similar, if not more similar, than concepts that are actually much closer together in the taxonomy. For instance, the induction of the concept *{wisdom_deity}* reveals *{Athena, Athene}* to be more similar to *{Minerva}*, a non-sibling, than it is to *{Aphrodite}*, a direct sibling. Though our approach to induction is not couched in terms of logic programming, as it is in [3, 12], we nonetheless believe it obeys the inductive philosophy of these approaches. In doing so, dynamic type creation also embodies the general philosophy of creativity espoused by John McCarthy in [10], which requires that for a solution to be creative, it must recruit or create concepts not directly mentioned in the problem specification. McCarthy’s viewpoint is interesting because it forces us to evaluate creativity not just on the utility of the end-product, which might equally be produced by the most banal of exhaustive searches, but on the selective means through which this end was achieved. As we have stated, and as our empirical analysis will support, dynamic types are created selectively, not as an end-product in themselves, but as a means to an end in generating insightful observations about ontological structure.

5 Empirical Analysis

In a DTH, new types are created opportunistically, in the context of specific metaphor interpretation or analogical reasoning tasks. For example, the types *{Hebrew_letter}* and *{Greek_letter}* are created in response to the analogy "What is the Jewish delta?". However, to enable us to test the coverage and scalability of our WordNet-based DTH, we have applied the type creation process to every one of the 69,780 unique noun sense entries of WordNet 1.6.

The strict limits on term reification imposed by definitions 1 and 2 mean that of the available 40,000 or so contents words in these noun glosses, only 2806 terms are actually reified, to create 9822 new dynamic types like *{cheese_dish}* and *{sea_deity}*. These dynamic types serve as hyponymic signposts for 2737 existing hypernyms, such as *{dish}* and *{deity}*, which allows these concepts to serve as the pivot types of future analogies. Under these 9822 signposts are indexed the 18,922 hyponyms (representing 27% of the WordNet noun taxonomy) from whose glosses the signposts were reified. Many of these hyponyms are connected to more than one signpost, via a collective total of 28,998 new hypernym-links. Each new type thus serves to unite an average of 3 existing concepts apiece.

One can think of the potential pairings in an ontology and the hierarchical similarity that they share as forming an abstract similarity space, in which similar concepts are positioned closer together than dissimilar concepts. The addition of new types to an ontology causes this similarity space to become warped in interesting ways. By introducing a finer level of taxonomic differentiation, certain concept pairs that did not previously share an immediate common parent are brought closer together because of their analogical potential, while other concepts pairs that do already share a common parent can, relatively speaking, be moved further apart if their common parent is less specific than a newly created dynamic

type. Our experiments show the effects of dynamic typing on WordNet to be widespread: sports are differentiated into team sports, ball sports, court sports, racket sports and net sports; constellations are divided according to whether they can be seen in the "northern" or "southern" hemispheres; food dishes are differentiated according to their ingredients, into cheese dishes, meat dishes, chicken dishes, rice dishes, etc.; letters are differentiated both by culture, giving Greek letters and Hebrew letters, and by relative position, so that "alpha" is both a *{1st_letter}* and a *{Greek_letter}*, while "Aleph" becomes both a *{1st_letter}* and a *{Hebrew_letter}*; and deities are differentiated to yield *{war_deity}, {love_deity}, {wine_deity}, {sea_deity}, {thunder_deity}, {fertility_deity}*, and so on. Overall, the most reified gloss term is "Mexico", which serves to signpost 34 different pivots (such as *{dish}*, grouping together *{taco}, {burrito}* and *{refried_beans}*), while the most differentiated pivot is *{herb, herbaceous_plant}*, which gains signposts to 134 sub-types like *{prickly_herb}*.

Table 1 presents a cross-section of the various sub-domains of *{deity, god}* in WordNet as they are organized by dynamic types such as *{supreme_deity}*. Where a mapping was not possible because a culture those not make a given distinction, N/A is used to fill the corresponding cell. In two cases, marked by (*), an adequate mapping was not be generated when one was culturally plausible; in the case of *{Odin}*, this is due to the gloss provided by WordNet 1.6, which defines Odin as a "ruler of the Aesir" rather than the "supreme" deity of his pantheon; in the case of *{Apollo}*, a Greco-Roman deity, the failure is due to this entity being defined solely as a Greek deity in WordNet 1.6.

Dynamic types primarily increase the precision, rather than the recall rate, of analogical mapping. Consider again the alphabet mapping task, in which the 24 letter concepts of the Greek alphabet are mapped onto the 23 letter concepts of the Hebrew alphabet, and vice versa. The recall rate for the Hebrew to Greek letter task, using both the DTH described here and the standard WordNet taxonomy, which we denote as the dynamic and static hierarchies respectively, is 100%. For the reverse task, mapping Greek to Hebrew letter concepts, the recall rate is 96% since Greek contains an extra letter than cannot be mapped to an equivalent Hebrew letter. However, the precision of the static WordNet hierarchy on these tasks is only 4%, since the static WordNet lacks the differentiating structure to pick out precise target mappings (see Figure 2). Consequently, every letter in the target alphabet is seen as an equally valid mapping. In contrast, the dynamic WordNet has sufficient differentiation to lock onto a precise target when one is available (see Figure 3). The precision of the dynamic DTH is 96% (Greek to Hebrew) and 100% (Hebrew to Greek).

Common Basis	*Greek*	*Roman*	*Hindu*	*Norse*	*Celtic*
supreme	Zeus	Jove	Varuna	Odin *	N/A
wisdom	Athena	Minerva	Ganesh	N/A	Brigit
beauty, love	Aphrodite	Venus	Kama	Freyja	Arianrhod
sea	Poseidon	Neptune	N/A	N/A	Ler
fertility	Dionysus	Ops	N/A	Freyr	Brigit
queen	Hera	Juno	Aditi	Hela	Ana
war	Ares	Mars	Skanda	Tyr	Morrigan
hearth	Hestia	Vesta	Agni	N/A	Brigit
moon	Artemis	Diana	Aditi	N/A	N/A
sun	Apollo	Apollo *	Rahu	N/A	Lug

Table 1. Mappings between sub-domains of *{deity, god}* in WordNet 1.6

The data of Table 1 allows for 20 different mapping tasks in the deities domain (Greek to Roman, Roman to Hindu, etc.). The average recall rate of the dynamic hierarchy is 61%, since some pantheons are less fleshed out than others. The Norse to Hindu mapping, for example, has a precision of just 30% for this reason. For the static hierarchy, average recall is significantly lower at 34%, since many concepts (such as Varuna and Aphrodite) are not indexed on the appropriate terms due to poorly defined glosses (e.g., Varuna is simply glossed as "supreme cosmic deity" in WordNet 1.6, and no explicit reference to Hinduism is made). Average precision for the dynamic hierarchy is 93.5%, with a loss of 6.5% precision

arising from the items marked (*) in Table 1. In contrast, average precision for the static hierarchy is just 11.5%, and would be lower still if incomplete glosses did not limit the number of incorrect answers that the static hierarchy approach can retrieve. That is, the problems that serve to reduce recall for the static hierarchy serve to artificially inflate its precision.

6 Conclusions

Analogy is a form of creative insight that recognizes the potential for two concepts that are separated in an ontology to be more similar than concepts that are taxonomically closer together. For instance, *{Ares}* and *{Mars}* are more similar than *{Ares}* and *{Zeus}*, despite the fact that the latter are taxonomic siblings while the former are not. Such insights, if achieved through automated means, can alleviate the structural problems that inevitably occur in manually constructed ontologies, especially those constructed on the ambitious scale of WordNet and Cyc which are naturally prone to problems of incompleteness and imbalance. The one-size-fits-all nature of such grand ontologies results in an organization that is often too undifferentiated for precise similarity judgments and too lopsided to support metaphor and analogical mapping.

A key symptom of these problems, and one that we can exploit, is the fact that English glosses or commentaries provide the ultimate level of differentiation in ontologies like WordNet and Cyc, so that one cannot truly differentiate two concepts without first understanding what their human-oriented glosses mean. In this paper we have described how new concepts, effectively dynamic ontological types, can be created by lifting implicit discriminators out of the flat text of these unstructured glosses and using them to construct new concepts. Our empirical results suggest that new type creation in a DTH constitutes a form of unsupervised learning about the conceptual dimensions that can most effectively organize a domain. For example, dynamic typing allows a DTH to learn for itself that deities are most commonly organized by dimensions such as War, Love, Fertility and Wisdom, and does so armed with no advance knowledge of the domain other that which can be found in WordNet.

Our experiments of section 5 were conducted using an information retrieval (IR) methodology, to compare the recall and precision of analogical mapping with and without the creation of dynamic types. IR offers a valuable perspective on analogical retrieval that suggests not only how we can evaluate our techniques empirically, but also how we can exploit them for practical ends. One avenue we are currently exploring is the development of an *analogical thesaurus* that will allow a user to search for variations on a concept that the user can only lexicalize indirectly, via metaphor and analogy [17, 18,19,20]. In contrast, a conventional thesaurus is predicated on an inherent circularity that only allows a user to search for variations of a concept the user already knows how to lexicalize directly. Suppose one wanted to find the Hindu, Roman or Semitic equivalents of the Greek gods Zeus, Ares and Athena, or to find the Muslim version of the bible, a church or a priest? Instead of the simple query "church" or "bible" that one would naturally use, to little effect, in a conventional thesaurus, one can pose cross-domain queries to an analogical thesaurus, such as "Muslim church" (returning mosque), "Hindu bible" (returning, amongst other candidates, the Vedas), "Celtic Ares" (returning Morrigan) and "Jewish German" (returning Yiddish).

The analogical thesaurus is just one application of a new branch of information retrieval we call *creative information retrieval* [21]. Conventional IR is an effective mechanism for text management that is extremely widespread, but one that can hardly be considered creative, at least not in the sense of [1,2, 4, 10, 23]. An IR system does not create new conceptual structures, or even reorganize existing ones, to pull in documents that describe, in novel and inventive ways, a user's information needs. Since language is a dynamic and highly creative medium of expression, the concepts that one seeks will therefore represent a

moving target for conventional IR systems. For instance, a creative IR system must predict that "comic books" can be creatively described as "graphic novels" if large parts of the relevant document space are not to be overlooked. Only by thinking creatively can an IR system effectively retrieve documents that express themselves creatively. This is the challenge of creative information retrieval: to imbue an IR engine with the conceptual tools and ontological fluidity needed to express itself creatively, so it can predict the creative ways in which a user's search concept might be communicated. We expect that research into creative metaphor and analogy will play a significant role in these exciting new enhancements to IR capabilities.

References

1. M. Boden, The Creative Mind (Abacus, London, 1990)
2. M. Boden, Creativity: A Framework for Research, Behavioural and Brain Sciences 17(3) (1994) 558-568.
3. S. Colton, Automated Theory Formation in Pure Mathematics, (Springer-Verlag, Berlin, 2002)
4. S. Colton, Creative Logic Programming, in: Proc. IJCAI'03 workshop on Creative Systems (Acapulco Mexico, 2003)
5. D. Fass, An Account of Coherence, Semantic Relations, Metonymy, and Lexical Ambiguity Resolution, in: S. I. Small, G. W. Cottrell and M.K. Tanenhaus, eds., Lexical Ambiguity Resolution: Perspectives from Psycholinguistics, Neuropsychology and Artificial Intelligence (Morgan Kaufman, San Mateo California, 1988)
6. S. M. Harabagiu, G. A. Miller and D. I. Moldovan, WordNet 2 - A Morphologically and Semantically Enhanced Resource, in: Proc. ACL SIGLEX Workshop on Standardizing Lexical Resources (Maryland, 1999)
7. J. Hutton, Aristotle's Poetics (Norton, New York, 1982)

8. G. Lakoff and M. Johnson, Metaphors we live by (University of Chicago Press, 1980)

9. D. Lenat and R. V. Guha, Building Large Knowledge-Based Systems (Addison Wesley, Reading Massachusetts, 1990)

10. J. McCarthy, Creative Solutions to Problems, in: Proc. AISB'99 symposium on AI and Scientific Creativity (Edinburgh Scotland, 1999)

11. G. A. Miller, WordNet: A Lexical Database for English, Communications of the ACM, 38(11) (1995)

12. S. Muggleton, Inductive Logic Programming, New Generation Computing 13 (1991) 245-286.

13. A. Ortony, The role of similarity in similes and metaphors, in: Ortony, A. ed., Metaphor and Thought (Cambridge University Press, Cambridge, U.K., 1979)

14. J. Pustejovsky, The generative lexicon, Computational Linguistics 17(4) (1991)

15 T. Veale and M. T. Keane, The Competence of Sub-Optimal Structure Mapping on Hard Analogies, in: Proc. IJCAI'97, the International Joint Conference on Artificial Intelligence, Nagoya, Japan (Morgan Kaufman, San Mateo California, 1997)

16. T. Veale and D. O'Donoghue, Computation and Blending, Cognitive Linguistics 11(3/4), 253-281.

17. T. Veale, Systematicity and the Lexicon in Creative Metaphor, in: Proc. ACL 2003 workshop on Figurative Language, (Sapporo, Japan, 2003)

18. T. Veale, Qualia Extraction and Creative Metaphor in WordNet, in: Proc. IJCAI'03 workshop on Creative Systems (Acapulco Mexico, 2003)

19. T. Veale, Dynamic Type Creation in Metaphor Interpretation and Analogical Reasoning, in: Conceptual structures for Knowledge Creation and Communication, LNAI 2746 (Springer, Berlin, 2003)

20. T. Veale, The Analogical Thesaurus: An Emerging Application at the Juncture of Lexical Metaphor and Information Retrieval, in: Proc. IAAI'03, (The AAAI Press, Menlo Park California, 2003)

21. T. Veale, Creative Information Retrieval, in: Proc. CICLing'04, (Springer Verlag LNCS, Berlin, 2004)

22. E. C. Way, Knowledge Representation and Metaphor, Studies in Cognitive systems (Kluwer Academic Publishers, Amsterdam, 1991)

23. G. Wiggins, Categorizing Creative Systems, in: Proc. IJCAI'03 workshop on Creative Systems (Acapulco Mexico, 2003)

Dynamic Creation of Analogically-Motivated Terms and Categories in Lexical Ontologies

Tony Veale,
School of Computer Science and Informatics, University College Dublin, Ireland.

1. Introduction

Ontologies are, for the most part, static organizations of categories and relations that attempt to model some aspect of the world. But such organizations can be strained to the breaking point when creative actions in this sub-world necessitate a dynamic response in the corresponding ontological structures. An important case in point is when, through some purposeful activity, an agent dynamically creates a new category of entity on the fly that must then be accommodated within the ontology (e.g., see Way, 1991; Veale, 2003a, 2005). This contrasts with the creation of new senses for existing ontological terms (as, for example, described in Pustejovsky, 1991). Such "ad-hoc" categories, as described by Barsalou (1983), are typically created in response to a specific goal or task, and may thus be considered task-specific or goal-specific constructs. Examples of ad-hoc categories include "things to take on a camping trip", "useful wedding presents", "substances from which sugar can be extracted", "objects that would make unusual murder weapons", and so on. Ad-hoc categories do not correspond to the existing hierarchical categories of an ontology, and their members are rarely clustered in the same localized area of the ontology. Rather, the members of a given ad-hoc category may be drawn from many different established categories. Ad-hoc categories thus constitute a horizontal rather than vertical slice of an ontology, cutting across conventional hierarchical structures, and as such, the creation of ad-hoc categories may intuitively be seen as a form of "lateral thinking" (de Bono, 1994).

The horizontal nature of ad-hoc categories means that they cannot easily be lexicalized with any of the labels associated with existing hierarchical categories; in fact, as seen above, the lexical label given to an ad-hoc category can be quite a mouthful, as a new multi-word expression must be constructed to capture the functional rather than taxonomic nature of the category. Nonetheless, some ad-hoc categories can be given compact labels that may subsequently find use as meaningful collocations in their own right and which merit their own individual listings in the lexicon. The creativity inherent in the construction of ad-hoc categories can thus apply at two different levels of representation, involving the creation of not just new ontological categories but of new lexical entries as well.

Of course, these levels of representation become one and the same when the ontology in question is a lexical ontology (see Veale *et al.* 2004, Hayes and Veale, 2005), that is, an ontology that concerns itself only with those units of meaning, called lexical concepts, that correspond to specific words or compound terms in a language. In this case, the creation of a new ontological category will correspond to the creation of a new lexical concept, forcing all new categories to assume a lexical label that serves a useful indexing role in the lexicon. For example, a new category like French-Food will serve to cluster together the various foodstuffs, dishes and wines that can be considered French into a single ontological category. At present, a number of large-scale lexical ontologies are available to support research in this area, among which WordNet a comprehensive electronic thesaurus of English whose design reflects psycholinguistic insights into the structure of the mental lexicon; nouns, verbs and adjectives are represented in different ways, and each lexical concept partakes in one or more relations to other concepts (such as hypernymy for nouns and troponymy for verbs). HowNet (see Dong 1998; Carpuat *et al.* 2002) is an equivalent Chinese ontology (with English translations) in which each lexical concept is associated with a propositional semantic structure. Both WordNet and HowNet are more properly described as *weak* ontologies since they exhibit neither the relational richness nor formal precision of the structures normally called ontologies by philosophers. In contrast, the Cyc ontology of Lenat and Guha (1991) is a rich axiomatic system explicitly designed to support logical inference and knowledge-based problem-solving, but which also maps word-forms onto a selection of underlying logical forms, both atomic and formulaic. However, though Cyc deserves to be called a strong ontology, it is not a lexical ontology, since natural language is not a motivating factor in its design; consequently, Cyc contains a high proportion of unlexicalized concepts and new concepts are not required to have corresponding lexical forms. Philosophers of language often contrast the role of the dictionary and the encyclopaedia when considering the knowledge demands of language; a lexical ontology aims to capture key aspects of both the dictionary and the encyclopaedia, and so constitutes the ideal framework in which to explore the mechanics of term creation.

Since ad-hoc categories are goal-specific, different goal contexts might give rise to different kinds of ad-hoc categories, suggesting that ad-hoc categories are best studied from the perspective of a specific cognitive goal or task. Now, one important task to which lexical ontologies have been directed is the construction and interpretation of lexical analogies. Analogy has been identified as a reasoning mechanism at several different levels of linguistic operation, from native speaker intuitions about pronunciation (see Baron, 1977) to intuitions about morphological inflection (see Trask, 1996) to intuitions about semantic relatedness (see Rumelhart and Abrahamson, 1973). At a lexico-conceptual level, analogies such as *Fructose is to Fruit as Lactose is to Milk* (e.g., see Veale 2003, 2004, 2005) exhibit creativity not only in their production, since they constitute novel linguistic artefacts, but also in their interpretation, which frequently requires the dynamic construction of new ad-hoc

categories. Consider the joke given in Freud (1905), and analysed as an analogy in Attardo *et al.* (2002):

"A wife is like an umbrella. Sometimes one takes a cab"

Attardo *et al*. (2002) provide the missing concept, Prostitute, to complete the analogy:

wife : prostitute :: umbrella : cab

To understand the analogy (and thus the joke) the listener must recognize that wives are *personal* lovers, while prostitutes are *hired* lovers; and that umbrellas are *personal* resources, while cabs are *hired* resources. This recognition necessitates the creation of the ad-hoc categories Personal-Lover and Personal-Resource (both sub-types of Personal-Belonging, to take a Victorian view of marriage and wives), as well as Personal-Resource and Hired-Resource. The burden of creativity is not borne solely by the creator of the joke, as the listener must also carry much of this burden through the creation of new categories that mirror the mind-set of the joker. Now, categories like Hired-Lover are highly goal-specific, and may not persist beyond the immediate context of the analogy that gives rise to them. However, if an ad-hoc category demonstrates some long-lasting value, its lexical label may also persist, to the point that it becomes a unit of common currency in the language and a permanent entry in a speaker's lexical ontology.

In this paper we consider how a particular kind of lexicalized ad-hoc category is created in a lexical ontology when the motivating task is that of analogical reasoning. We argue that the interpretation of lexical analogies often necessitates the creation of new conceptual categories that in turn necessitate the creation of new lexical items. These lexical terms may be as short-lived as the analogies themselves, but a corpus analysis can be used to reveal those terms that have sufficient durability to merit a place in the lexical ontology. For our current purposes, we ground our investigation in the context of WordNet, and explore a variety of ways in which analogy can be used to drive the creation of lexical innovations that do not already exist within the WordNet lexicon.

We are careful to note that, in the context of a lexical ontology, the term “analogical categorization” is an ambiguous one. It can mean either the creation of new categories, like Hired-Lover and Personal-Resource, to resolve a particular analogy, or it can denote the use of analogy as an explicit term-creation mechanism, in much the same way that analogy can be used to suggest spelling, morphology and pronunciation. Since each reading denotes a process that is meaningfully performed within a lexical ontology, the ambiguity is a benign one, reflecting the multipurpose nature of lexical analogy. For the sake of completeness, we consider both of the foregoing uses in this paper. In sections two and three we explore the use of analogy as an explicit and quite deliberate mechanism of term creation, while in section four we consider how new ad-hoc concepts and their corresponding lexicalizations can arise as by-products from the interpretation of lexical analogies. More specifically, section two considers how analogy can be used to produce new words, while sections three and four explore the role of analogy in the creation of new compound terms. In section five we consider the application of these ideas to the WordNet lexical ontology, which allows us to empirically evaluate their effectiveness in the context of an analogical retrieval task. We then conclude with some remarks on the limitations of these ideas in section six.

2. Analogy as a Mechanism of Word Creation

To begin, one might well consider how analogy is implicated in the determination of morphological inflections, even in the face of stronger and more accepted linguistic principles. Indeed, as noted by Trask (1996), analogies can be so persuasive that any

fallacious conclusions that can be drawn may seem even more natural than those of a first principles analysis. For instance, the lexical analogy in (1) is compelling even though the *+us→+i* pluralization rule is valid only for words of Latin origin, as exemplified by radius/radii and succubus/succubi.

(1) cactus : cacti :: octopus : octopi

The correct inflection, following the Greek origins of the word octopus, is *octopodes*, yet this is far less favoured by English speakers, to the extent that the automatic spelling corrector provided by Microsoft Word deems “octopi” to be valid and “octopodes” to be a misspelling. This usefulness of analogy in dealing with irregular plurals even extends to the treatment of regular verbs, where a compelling analogy can make even a regular verb like “dive” seem irregular. The analogy *drive:drove:dive:dove* will rightly strike some readers as invalid, yet many Eastern American speakers strongly prefer *dove* to *dived* (Trask, ibid). In contrast, the analogy *teach:taught::catch:caught* seems a valid one, though as Trask notes, *catched* is actually the historically favoured past tense of catch. Again, however, analogy prevails to the extent that most spelling checkers will flag “catched” as a misspelling (no doubt due to the fact that spelling checkers are based on a corpus analysis of how language is actually used, rather than how it should be used).

Morpheme-level analogies can do more than suggest inflection patterns, and can even be used to derive new words and meanings that frequently exhibit a high degree of lexical creativity. Consider an example of analogy-based derivational morphology:

(2) astronomy : astronomer :: gastronomy : gastronomer

Neither WordNet nor the spelling checker for Microsoft Word recognize “gastronomer” as a valid word, though a web-search reveals that it is a real word with the semantics that one would expect from the analogy (i.e., a specialist in gastronomy). The analogy in (2) is semantically sound since Astronomy and Astronomer are strongly related concepts, but as one allows the analogy to veer towards the speculative, and to rely as much on sound similarity as semantic similarity, one can achieve even more innovative results, as in (3).

(3) astronomy : astronaut :: gastronomy : gastronaut

The relation between astronomy and astronaut is a good deal more tenuous than that between astronomy and astronomer, but a relation does exist (one *observes* the stars, the other *explores* the stars, in name at least). Indeed, one can argue that Astronaut is itself analogically derived from Argonaut. A gastronaut might thus denote anything from an adventurous gastronomer to a food tourist; at the very least, we know that a gastronaut is a person, with some of the signal characteristics of an astronaut (bravery, perhaps), who takes his directions from the field of gastronomy.

If it seems that the process of lexical creativity can be as strongly influenced by phonetic concerns as semantic concerns, this should not be too surprising a conclusion. New words survive and thrive for a whole host of reasons, but an important factor in their survival is euphony: natural sounding words are more likely to secure a lasting place in the lexicon than those that are difficult to pronounce. Analogies with existing words can transplant the euphony of an original form onto a newly minted neologism only if phonetic similarity is also allowed to influence the mapping. Indeed, the most innovative creations may give so much prominence to phonetic similarity that an analogy may lack a credible semantic basis. Nonetheless, as demonstrated in (4), terms predicated on a false analogy can still be seen as lexically innovative:

(4) astronomy : astrodome :: gastronomy : gastrodome

Of course, there is no real semantic connection between astronomy and astrodome. Nonetheless, the word "gastrodome" can be seen as a deliberate malapropism that

amply suggests a place where gastronomy is performed, and perhaps even celebrated[1]. The analogy works, despite its lack of semantic grounding, because "dome" is itself a word denoting a large enclosed space where people congregate, like an arena or stadium. A second, implied analogy can be used to tease out its precise meaning:

(5) astrodome : stadium :: gastrodome : restaurant

That is, just as an astrodome is a large, impressive stadium, a gastrodome is a large impressive restaurant (where restaurant is itself implied by the morphological conjunction of gastronomy and place in "gastrodome"). Ultimately, "Gastrodome" is preferable to the neologisms "Gastroarena" and "Gastrostadium" in part because the largely phonetic analogy ensures that it is a euphonious combination of morphemes, each of which can be considered in isolation to provide a compositional meaning to the neologism as a whole.

3. Analogy as a Mechanism of Compound-Term Creation

Moving from the level of morphemes to that of words, analogy again reveals itself as a powerful force in the creation of compound terms. Consider the analogy of (6):

(6) Greek-Alphabet : Hebrew-Alphabet :: Greek-Deity : Hebrew-Deity

The analogy captures a basic symmetry both in the way concepts can be differentiated and how such differentiations are lexically expressed as compound terms. Both "Greek" and "Hebrew" denote a cultural amalgam of people, language and belief, so it makes sense to conclude that if "Greek" can be used to culturally differentiate a particular concept, then so can "Hebrew". In fact, WordNet contains only three of the four compound terms in the above analogy: "Greek-Alphabet", "Hebrew-Alphabet"

[1] As of August 2005, the Google search engine retrieves almost 7000 documents for "gastronomer", a surprising 70,000 for "gastronaut", a mere 7 for "gastronautics" and over 700 for "gastrodome".

and "Greek-Deity". The lexical concept "Hebrew-Deity" is not listed as a WordNet entry, most likely because it is deemed to have little indexing value; while there are many Greek deities listed in WordNet that would structurally benefit from the clustering offered by the hypernym Greek-Deity, only one deity, Jehovah, is listed as having a Hebrew origin. Yet the concept seems logically well-formed, and a usage analysis (using the World Wide Web as a corpus) reveals that the term "Hebrew deity" has relatively widespread acceptance[2]. The analogy suggests then that, on the basis of the similarity between Greek and Hebrew, WordNet should incorporate the lexical concept Hebrew-Deity. Were it to do so, its treatment of deities would become more systematic and balanced, with each proper deity (such as "Zeus", "Mars" and "Jehovah") instantiating a compound category that denoted its cultural basis.

Nonetheless, a simple proportional analogy like (6) may seem a weak basis on which to predict the existence of a new term. Consider the analogies of (7) and (8):

(7) Roman-Alphabet : Greek-Alphabet :: Roman-Empire : Greek-Empire

(8) Roman-Alphabet : Hebrew-Alphabet :: Roman-Empire : Hebrew-Empire

Here we seem to be predicating the possession of an empire on the existence of a unique alphabet, but an alphabet alone does not an empire make. The analogy of (7) holds true, since there is a historical entity called the "Greek Empire", but the term "Hebrew Empire" can only be used metaphorically, perhaps to refer to the Jewish diaspora. However, this is not to say that analogy cannot be of use here, for what (7) and (8) fail to reveal is the variety of different analogies that support (7), and the comparative dearth of analogies that support (8). In addition to alphabets, the Greeks and the Romans both possessed their own mythologies, architectures, religions and

[2] Google returned 907 documents containing the term "Hebrew deity" in August 2005. Most occurred in the context of a definite article, suggesting that this category of deity is mostly conceived of as a definite description for a unique entity.

deities. This semantic isomorphism suggests that if Rome possessed its own empire, it is at least meaningful to consider the possibility of a Greek empire also. The analogy in (7) is therefore strengthened by the lexico-conceptual fit between the concepts Roman and Greek, while the analogy in (8) is much weaker because of the lack of a coherent fit. This "fit" is not a measure of ontological closeness, but a measure of the overlap between the set of affordances possessed by both concepts. We can loosely estimate this set of affordances by observing the lexical behaviour of each term and how it relates to others. These observations will require us to define a set of basic term composition and decomposition operators along the following lines:

$\mathbf{U_M}$ **{X}** : *Usage as modifier*: return a set of all compound terms such that the modifier of each is a member of the set {X}.

E.g., $\mathbf{U_M}${Greek, Roman} = {Roman-deity, Roman-Empire, Greek-deity, ...}

$\mathbf{U_H}$ **{X}** : *Usage as head*: return a set of all compounds such that the head of each is a member of the set {X}.

E.g., $\mathbf{U_H}${Greek, Roman} = {Ancient-Greek, Modern-Greek, Times-Roman, ...}

M {X} : *get modifiers*: return the set of all modifiers of all compounds in {X}

E.g., **M**{Ancient-Greek, Times-Roman, ...} = {Ancient, Times, ...}

H{X} : *get heads*: return the set of all heads of all compounds in {X}

E.g., **H**{Ancient-Greek, Modern-Greek, Times-Roman, ...} = {Greek, Roman, ...}

C ({X}, {Y}) : *combination*: return the set of all possible compound terms whose modifier is in {X} and whose head is in {Y}

E.g., **C**({Greek, Hebrew}, {Alphabet, Deity}) = {Greek-Alphabet, Greek-Deity, …}

These operators allow us to dissect existing compound terms into their component parts (modifier *and* head), retrieve compound terms with a particular sub-component (modifier *or* head), and create novel combinations of these sub-components (modifiers *crossed with* heads). We can thus estimate the fitness of a novel compound X-Y in terms of the set of known compounds that support it, as follows (where ***L*** here denotes the set of all lexical items in the lexicon, i.e., all known terms):

$$\textit{support-set}(\text{X-Y}) = \mathbf{C}(\mathbf{M}(\mathbf{U_H}\{Y\}), \mathbf{H}(\mathbf{U_M}\{X\})) \cap \boldsymbol{L}$$

For example, consider the support set for the novel compound Hebrew-Deity from (6):

$$\begin{aligned}\textit{support-set}(\text{Hebrew-Deity}) &= \text{C}(\mathbf{M}(\{\text{Greek-Deity, Roman-Deity, Semitic-Deity, ...}\}), \\ &\quad \mathbf{H}(\{\text{Hebrew-Alphabet, Hebrew-Calendar, …}\})) \cap \boldsymbol{L} \\ &= \mathbf{C}(\{\text{Greek, Roman, Semitic, Hindu, Celtic, Norse, …}\}, \\ &\quad \{\text{Alphabet, Calendar, Lesson}\}) \cap \boldsymbol{L} \\ &= \{\text{Greek-Alphabet, Roman-Alphabet,} \\ &\quad \text{Roman-Calendar, Hindu-Calendar}\}\end{aligned}$$

In effect, this set of four existing compounds represents the lexico-conceptual *cross-product* of the lexical concepts Hebrew and Deity. The larger the cross-product, the greater the potential interaction – and the greater the lexical fit – between both terms. One might think it strange that a lexical concept like Hindu-Calendar should support a term like Hebrew-Deity, but the intuition at work here is that deities and calendars

appear to be differentiated in the same kind of way (e.g., culturally) and thus possess many of the same affordances.

Working backward from this formulation of fitness, we can formulate a generation mechanism for producing new compound terms from old, one that implicitly incorporates the notion of lexical analogy. Consider that when a compound like Hebrew-Deity is generated from an analogy involving Hebrew-Alphabet, the head term is effectively *modulated* from Alphabet to Deity (see Veale *et al.* 2004; Hayes and Veale, 2005). Head modulation thus offers an alternate perspective on the generation process. Imagine that $new_{\mathbf{H}}$(X-Y) is a function that derives, via head modulation, a set of novel compound terms from an existing term X-Y. Using the operators above, $new_{\mathbf{H}}$ can be formulated as follows:

$$new_{\mathbf{H}}(\text{X-Y}) = \mathbf{C}(\{X\}, \mathbf{H}(\mathbf{U_M}(\mathbf{M}(\mathbf{U_H}\{Y\}))\backslash\{Y\})) \setminus \boldsymbol{L}$$

For example,

$$\begin{aligned}new_{\mathbf{H}}(\text{Muslim-Calendar}) &= \mathbf{C}(\{\text{Muslim}\}, \mathbf{H}(\mathbf{U_M}(\mathbf{M}\{\text{Hebrew-Calendar,} \\ &\qquad \text{Roman-Calendar,} \\ &\qquad \text{Hindu-Calendar ...}\}))) \setminus \boldsymbol{L} \\ &= \mathbf{C}(\{\text{Muslim}\}, \mathbf{H}(\mathbf{U_M}\{\text{Hebrew, Roman, Hindu, ...}\})) \setminus \boldsymbol{L} \\ &= \mathbf{C}(\{\text{Muslim}\}, \\ &\quad \mathbf{H}(\{\text{Hebrew-Alphabet, Roman-Deity, ...}\})) \setminus \boldsymbol{L} \\ &= \mathbf{C}(\{\text{Muslim}\}, \{\text{Alphabet, Deity, Empire, ...}\}) \setminus \boldsymbol{L} \\ &= \{\text{Muslim-Alphabet, Muslim-Deity, Muslim-Empire, ...}\}\end{aligned}$$

The resulting set of speculative compounds must now be evaluated for lexico-conceptual fitness, using the *support-set* measure described earlier. At this point we expect those compounds with the greatest fit to be the best candidates for lexical innovation and subsequent admission to the lexicon. Before taking this final step, which could potentially corrupt the lexicon, we can apply a further fitness filter by demanding that each new term present a given number of times in a given corpus (such as the WWW). Veale *et al.* (2004) report experimental findings which suggest that the probability of finding a newly generated term in a corpus such as the WWW increases with the size of the support set for that term. The larger the support set, then, the safer it is to conclude that a lexical innovation is in fact meaningful.

If new compounds can be generated by modulating the head component of existing terms, it follows that generation can also proceed via a process of modifier modulation, whereby the modifier component of an existing term is modulated according to an implicit analogy. We can formulate modifier modulation as follows:

$$new_{\mathbf{M}}\text{(X-Y)} \quad = \mathbf{C}(\mathbf{M}(\mathbf{U_H}\ (\mathbf{H}(\mathbf{U_M}\{X\})))\backslash\{X\},\ \{Y\})\ \backslash\ \boldsymbol{L}$$

For example,

$$\begin{aligned} new_{\mathbf{M}}\text{(Hebrew-Alphabet)} &= \mathbf{C}(\mathbf{M}(\mathbf{U_H}(\mathbf{H}(\{\text{Hebrew-Lesson, Hebrew-Calendar}\})), \\ &\quad \{\text{Alphabet}\})\ \backslash\ \boldsymbol{L} \\ &= \mathbf{C}(\mathbf{M}(\mathbf{U_H}\ (\{\text{Lesson, Calendar}\}),\ \{\text{Alphabet}\})\ \backslash\ \boldsymbol{L} \\ &= \mathbf{C}(\mathbf{M}(\{\text{German-Lesson, …, Muslim-Calendar…}\}), \\ &\quad \{\text{Alphabet}\})\ \backslash\ \boldsymbol{L} \\ &= \mathbf{C}(\{\text{German, French, Muslim, …}\},\ \{\text{Alphabet}\})\ \backslash\ \boldsymbol{L} \\ &= \{\text{German-Alphabet}^{3}, …, \text{Muslim-Alphabet}^{4},...\} \end{aligned}$$

So we speculatively create the compounds German-Alphabet because of an implicit analogy between Hebrew-Lesson and German-Lesson, and Muslim-Alphabet because

[3] "German Alphabet" retrieved 6640 documents from the WWW via Google in August 2005. The compound is meaningful because German contains characters (like ß) and diacritics that do not appear in the Roman alphabet.

[4] "Muslim Alphabet" retrieved a mere 28 documents on the same date, which for the most part assume it to mean "Arabic Alphabet". In contrast, the compound term "Muslim Empire" was found in over 28,000 documents.

of the implicit analogy between Hebrew-Calendar and Muslim-Calendar.

Alternately, we can conceive of a generation scheme in which both the modifier and head of an existing term can be modulated simultaneously, as formulated below:

$$new_{\mathbf{MH}}\text{(X-Y)} \quad = \mathbf{C}(\mathbf{M}(\mathbf{U_H}\ \{Y\})/\{X\},\ \mathbf{H}(\mathbf{U_M}\ \{X\})/\{Y\})\ \backslash\ \boldsymbol{L}$$

This highly speculative formulation generates the cross-product of all modifiers that can apply to Y with all heads that can be modified by X. Again, many untenable combinations will be produced, but following Veale (2004), we can expect that those with a sizeable support set will be meaningful.

3.1. Experimental Support

This belief is further supported by an experiment in which 100,000 novel compounds were chosen at random from the set of all compounds that can be created via the modulation of existing WordNet compounds. These new compounds, created using the formulations of $new_{\mathbf{MH}}$, $new_{\mathbf{M}}$ and $new_{\mathbf{H}}$ given above, are grouped into different categories according to the size of their support sets; for example, compounds with a support set of 5 other compounds are organized under the category *group-5*, and so on. We can thus estimate the probability that a compound with a support set of size *n* will be validated on the WWW as that fraction of those elements of *group-n* that are so validated. In fact, there is a significant positive correlation (0.4) between *n* and the probability that an element of *group-n* will be validated via web-search. This correlation remains stable whether modifier modulation ($new_{\mathbf{M}}$), head modulation ($new_{\mathbf{H}}$) or simultaneous modulation ($new_{\mathbf{MH}}$) is used to generate the test data.

3.2. Phonetic Analogies Revisited

We are now in a better position to consider the phonetically-inspired morpheme-level analogies of section 2. In each case we can view a multi-morphemic word as a compound term, composed of morphemes rather than words, such that one of these morphemes is modulated by an implicit analogy. Consider again the analogy of (2):

(2′) astro-onomy : astro-onomer :: gastro-onomy : gastro-onomer

Now we can view "gastronomer" as a product of head modulation, where the head morpheme "-onomy" is transformed into the morpheme "-onomer" on the basis of an analogy with astronomy:astronomer. A similar process occurs in (3′), with the additional phonetic similarity between "astro-" and "gastro-" ensuring that the modulation, which is morphologically sound, also produces a euphonious result.

(3′) astro-onomy : astro-onaut :: gastro-onomy : gastro-onaut

Note that the modulation perspective saves us from having to rationalize a relationship between astronomy and astronaut, allowing us instead to view them as words that share a common modifier "astro-". This in turn allows us to exploit analogies like that of (4′) where no such semantic relationship exists:

(4′) astro-onomy : astro-dome :: gastro-onomy : gastro-dome

The terms "gastronaut" and "gastrodome" each have a singleton support set, corresponding to a single analogy of (3′) and (4′) respectively. In lieu of substantial support, however, these analogues are grounded by a phonetic similarity to their supports and this provides the requisite credibility for the new terms. That is, the similarity between "gastro-" and "astro-" is itself a support for the new terms. Note also that order is important, yielding a bootstrapping effect as new terms are incrementally accepted into the lexicon. For instance, if (4′) is processed after (2′) and (3′), the support for "gastrodome" can be determined as follows.

support-set(gastro-dome) = **C**(**M**({astro-dome}),

H({gastro-onomy, gastro-onomer, …})) ∩ ***L***

= **C**({astro-}, {-onomy, -onomer, -onaut}) ∩ ***L***

= {astronomy, astronomer, astronaut}

4. Ad-Hoc Categories as Analogical By-Products

Thus far we have considered the deliberate and explicit creation of new terms and categories whose existence is predicated on an implicit analogy (i.e., where the analogy is implicit in the workings of $new_{\mathbf{M}}$, $new_{\mathbf{H}}$ and $new_{\mathbf{MH}}$). We now consider the situation where a new lexical concept is created implicitly, as a by-product of the interpretation of an explicit analogy. For instance, consider the analogies in (9):

(9a) Zeus : Greek :: Jupiter : Roman

(9b) Zeus : Greek :: ??? : Roman

(9c) "Zeus is the Greek Jupiter"

The analogy of (9a) establishes an explicit mapping between Zeus and Jupiter and between Greek and Roman, suggesting that Zeus is the Greek equivalent of Jupiter. The variant in (9b) employs an elliptical form common to test questions such as the S.A.T., and requires us to provide the missing information; in effect, it equates to the question "What or Who is the Roman Zeus"? In contrast, the variant of (9c) assumes a compressed natural language form that can also be considered a metaphoric expression (e.g., see Hutton, 1982).

The implicit relation common to (9a), (9b) and (9c) appears to be "deity of": Zeus is a deity of the Greeks, while Jupiter is a deity of the Romans. However, consider the longer form of this analogy in (10):

(10) Zeus is to Greek as

a. *Skanda is to Hindu*

b. *Thor is to Norse*

c. *Jupiter is to Roman*

d. *Brigit is to Celtic*

e. *Donar is to Teutonic*

Each of the candidate pairings in (10) can be seen as instantiating the "deity of" relationship, so a more specialized relationship is clearly at work here. In fact, the correct relationship is "supreme deity of", since this is the only conceptual relationship for the stem pairing that picks out just one of the five possible candidates. Now, WordNet contains the concept Deity, so one can imagine constructing the relationship "deity of" from this concept in a relatively straightforward fashion. But WordNet does not contain the concept Supreme-Deity, and for good reason: it is not a conventional collocation, and its meaning is simply a compositional function of existing terms. One of two situations must therefore hold: either the concept already exists but is not lexicalized; or else neither the concept nor its lexicalization exists prior to the analogy. In either case, we can reasonably assume that the lexical term "supreme deity" is constructed especially to resolve the analogy.

Not all such analogies require us to construct new lexical concepts. Consider the analogy in (11), which can be seen as a close conceptual neighbour of (9a):

(11) Ares : Greek :: Mars : Roman

Here it is the relationship "war god of" that connects Ares to Greek and Mars to Roman. In this case, however, WordNet does contain the concept War-God, while its lexicalization "war god" is such a conventional collocation that few would argue that it is constructed especially for the purpose of this analogy. However, this is not to say

that the interpretation of (11) should be substantially different from that of (9) or (10). We can still presuppose that for each analogy, the same process is employed to construct a relational category between each concept in each pairing. In the case of (11), this relational category (War-God) will correspond to an existing lexical concept, while in (9) and (10) it will result in a lexical innovation ("supreme deity") that may be added to the lexicon following an assessment of its support set or a corpus analysis.

The construction of these relational categories raises two key questions: first, where do the component parts such as "war", "supreme" and "deity" come from; and second, why are these components, rather than others, selected? The lexicon or lexical ontology presumably plays a central role in resolving these questions, which further begs the question of what theory of the lexicon we should adopt. To remain as agnostic as possible, let us assume a rather simple, feature-theoretic view of the lexicon. Let ***F*** denote a function that maps a lexical concept onto a set of component features. Furthermore, let us assume that these features can be of one of two types. Taxonomic features, denoted with a **↑**, are those that indicate the position of a concept in the lexical ontology. Associative features, denoted with a **@**, are those that predicate descriptive properties of the concept. For instance, consider Zeus again:

F*(Zeus) = {↑deity, @Greek, @supreme, @mythology, @Olympus}*

Thus, Zeus is a deity that is Greek and supreme, associated with both mythology and Olympus. In contrast, we can define Jupiter as follows:

F*(Jupiter) = {↑deity, @Roman, @supreme, @mythology, @rain}*

Jupiter is thus a deity that is Roman and supreme, associated with both mythology and rain (in the guise of *Jupiter Pluvius*). This feature-level decomposition suggests a means whereby new categorizations can be created for a given concept. Consider the following formulation of a function ***alt***, which derives a set of alternate categorizations for a concept by constructing alternate compositions of elements in ***F***:

$$\boldsymbol{alt}(A) \quad = \quad \boldsymbol{C}(\ \{X \mid \mathbf{U}_{\mathbf{M}}\{X\} \neq \{\} \wedge @X \in \boldsymbol{F}(A) \ \wedge\ \uparrow X \notin \boldsymbol{F}(A)\},$$
$$\{Y \mid\ \mathbf{U}_{\mathbf{H}}\{Y\} \neq \{\} \wedge \uparrow Y \in \boldsymbol{F}(A)\})$$

That is, the set of alternate categorizations of A comprises just those compound terms that can be created by combining the associative feat ures of A that have in the past been used as compound modifiers with the taxonomic features of A that have in the past been used as compound heads. The resulting compound terms are thus well-formed with respect to the lexicon and the language that it represents. Note that this formulation of ***alt*** prohibits the hypernymic terms of a concept (like ↑*deity* for Jupiter) from serving as a modifier in any alternative categorization of the concept, since this is a combination strategy rarely seen among English compounds[5].

Now, a simplistic view of analogy, based on the Aristotelian account (see Hutton, 1982), might attempt to reconcile Zeus and Jupiter by seeking a common taxonomic feature (e.g., ↑deity) in both representations, but as demonstrated by (10), a genus term alone lacks discriminatory power. We need a common category that combines the Aristotelian notions of both genus *and* differentia. Given an analogical pairing A:B, we can construct this category using the function **adhoc**, formulated as follows:

$$\boldsymbol{adhoc}(A{:}B) \quad = \{X\text{-}Y \ \mid\ X\text{-}Y \in \boldsymbol{alt}(A) \cap \boldsymbol{alt}(B)$$
$$\wedge\ \neg(\exists P\ \uparrow P \in \boldsymbol{F}(A) \wedge \uparrow P \in \boldsymbol{F}(B) \wedge \uparrow Y \in \boldsymbol{F}(P))\ \}$$

Expressed in English, ***adhoc***(A:B) generates a set of compound terms X-Y such that: i) X-Y is an alternative categorization of both A and B; and ii) there is no other shared

[5] Generally speaking, the modifier of a compound term denotes a property of the head (as in "wax paper") or a concept from which a property is transferred to the head (as in "beehive hairdo") or a concept to which the head relates via slot-filling (as in "harpoon gun"). In some hybrid compounds, both the modifier and the head denote a hypernym of the compound, as in "sofa bed", but these are so rare as to be safely precluded from the current analysis. Were we to allow hybrid compounds, the formalism given here would surely overgenerate.

taxonomic feature of A and B (P, say) that is more specific than Y.

An analogy *A:B::C:D* is well-formed if precisely the same relationship holds between A and B and between C and D. For example, the analogy *Zeus:Hindu::Jupiter:Roman* is malformed because Zeus is not Hindu but Greek. Thus:

$$\boldsymbol{wellformed}(A{:}B{::}C{:}D) = \quad (\exists M\ \ M\text{-}B \in \boldsymbol{alt}(A)\ \wedge\ M\text{-}D \in \boldsymbol{alt}(C))$$
$$\vee\ \ (\exists H\ \ B\text{-}H \in \boldsymbol{alt}(A)\ \wedge\ D\text{-}H \in \boldsymbol{alt}(C))$$
$$\vee\ \ (\exists M_1M_2H_1H_2\ \ M_1\text{-}H_1 \in \boldsymbol{alt}(A)\ \wedge\ M_1\text{-}H_2 \in \boldsymbol{alt}(C)$$
$$\wedge\ \ M_2\text{-}H_1 \in \boldsymbol{alt}(B) \wedge\ M_2\text{-}H_2 \in \boldsymbol{alt}(B))$$

The first disjunct covers the situations where B and D are super-ordinates of A and C (as in the analogy *ewe:sheep::hen:chicken* where coherence is given by the relations female-sheep and female-chicken). The second disjunct covers the situations where B and D are features of A and C (as in the analogy *Athena:Greek::Ganesh:Hindu*). The third disjunct, the most complex, covers those situations where B and D are in some sense antonyms of A and C (as in *wife:prostitute::umbrella:cab*). Now, well-formedness does not always imply solvability; for that, there must exist a relationship between A and C that is mirrored between B and D. Thus, given the analogy *A:B::C:D*, we additionally expect that it has a non-empty relational basis:

$$\boldsymbol{basis}(A{:}B{::}C{:}D) = \boldsymbol{adhoc}(A{:}C) \neq \{\}$$

That is, the pairing A:C in a proportional analogy should share at least one relational category if *A:B::C:D* is to be considered a solvable analogy. As formulated above, ***basis*** may return a set containing a plurality of categories. In the case of analogies like (9a) and (11), it is sufficient that this be a non-empty set. But in the case of long-form analogies like (10), where a stem pairing must be matched with just one other in a group of candidate pairings, it may be possible that multiple candidate pairings share a non-empty relational basis with the stem pairing. In this case, one must choose the candidate with the *strongest* relational basis. Since each element returned by ***basis*** is a conceptual category, we can determine the discrimination strength of each category by considering it from an extensional perspective. Given two categories in the relational basis of an analogical pairing, e.g., supreme-deity and Greek-deity, the strongest category is taken to be that which has the smallest extension (and which is thus the most discriminating). The extension of Greek-deity is larger that that of supreme-deity (108 members versus 6 members in WordNet), so we take supreme-deity to be the stronger category on which to ground an interpretation.

What of partial analogies like (9b), which form the basis of both examination questions (where a student must provide the missing information) and metaphoric allusions? In such cases, a suitable analogue must be retrieved to complete the analogy, using the available information as a retrieval cue. We can formulate a retrieval-oriented variant of ***basis*** as follows:

$$\textbf{\textit{basis}}(A{:}B{::}???{:}D) = \{X\text{-}Y \mid \exists C \ @D \in \textbf{\textit{F}}(C) \wedge X\text{-}Y \in \textbf{\textit{adhoc}}(A{:}C)\}$$

If the lexicon is sufficiently indexed, as one might expect in a structured lexical ontology, it should be relatively straightforward to identify C using D as an index.

5. Analogical Retrieval in WordNet

The comprehensive scale of WordNet as a lexical database of English word meanings, with over 100,000 lexical concepts, allows us to put the intuitions and formulations of previous sections to the test. The specific task we propose in this section is that of analogical retrieval (see Veale, 2003b; Veale, 2004): given a lexical concept in one domain, such as "Zeus", and a modifier that denotes another domain, such as "Roman",

we seek to retrieve those concepts in the modifier domain that are meaningful analogies for the original head concept. The retrieval task is thus a question-answering task, in which we attempt to find answers for queries such as "Who is the Norse Zeus?" and "Who is the Hindu Athena". For balance, we shall conduct our test in two different domains of knowledge, namely deities and alphabets. The deities domain is quite well represented in WordNet, while structurally, the alphabetic domain is relatively impoverished. We shall demonstrate that the creation of ad-hoc categories that are subsequently admitted to the lexicon can significantly improve the state of these impoverished domains.

We concentrate our efforts then on the noun section of WordNet, which contains over 70,000 taxonomically organized entries. In addition to this taxonomic information, WordNet associates a textual gloss with each entry, much like that offered by a regular dictionary. For example, WordNet associates the following information with the concepts Zeus, Jupiter, Alpha and Aleph:

Zeus:	*Taxonomy*	= {Greek-deity is-a deity, is-a god is-a …}
	Gloss	= "The supreme god of ancient mythology"
Jupiter:	*Taxonomy*	= {Roman-deity is-a deity, is-a god is-a …}
	Gloss	= "(Roman mythology) supreme god of Romans"
Alpha:	*Taxonomy*	= {letter is-a character is-a written-symbol is-a …}
	Gloss	= "the 1st letter of the Greek alphabet"
Aleph:	*Taxonomy*	= {letter is-a character is-a written-symbol is-a …}
	Gloss	= "the 1st letter of the Hebrew alphabet"

Unfortunately, WordNet does not offer an explicitly feature-theoretic description of each lexical concept, such as that provided by our function ***F***. However, we can approximate the corresponding ***F*** for WordNet by assuming that the textual gloss of each concept is, in fact, a bag of associative features; we simply eject any non-content words (such as determiners, prepositions, and so on), and merge the resulting word set with the set of taxonomic parents that is explicitly provided by the WordNet. Thus, from WordNet we derive the following mappings for ***F***:

F(Zeus): = {↑Greek-deity ↑deity @supreme @god @ancient @mythology}

F(Jupiter): = {↑Roman-deity ↑deity @Roman @supreme @god @Romans}

F(Alpha): = {↑letter ↑character @1^{st} @letter @Greek @alphabet}

F(Aleph): = {↑letter ↑character @1^{st} @letter @Hebrew @alphabet}

Applying the function ***adhoc*** to these representations, we obtain the following:

adhoc(Zeus:Jupiter) = {supreme-deity}

adhoc(Alpha:Aleph) = {1^{st}-letter, alphabet-letter}

Note that the ad-hoc concepts god-deity and letter-letter, though seemingly possible from the given values of ***F***, are not created because of the definition of ***alt*** as formulated earlier (i.e., no taxonyms as modifiers). Note also that ***adhoc*** returns two different categories for the pairing of Alpha with Aleph. In this case, based on the extension of both categories, 1^{st}-letter is deemed the stronger of the two. In fact, an extensional analysis reveals that the extension of 1^{st}-letter (with just two members) is a proper subset of that of alphabet-letter (with 49 members), which suggests that 1^{st}-letter is a specialization of the category alphabet-letter.

Figure 1 illustrates the taxonomic structure of the letter domain in WordNet before any letter analogies (of the form *Alpha:Greek::???:Hebrew*) have been interpreted. Note the general paucity of organizational structure here: each letter from each alphabet is forced to share the same super-ordinate category, letter, and no attempt is made to gather letters from different alphabets under separate super-ordinates.

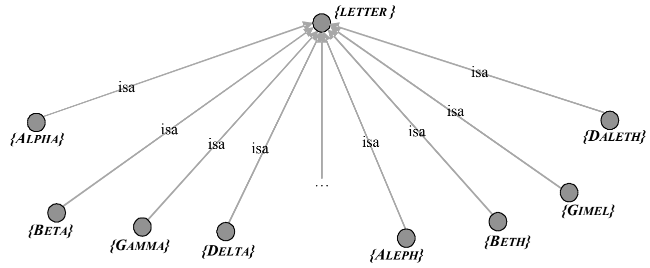

Figure 1: *The structure of the Greek and Hebrew letters domain in WordNet*

This picture changes dramatically once each letter in the Greek alphabet is placed in analogical alignment with its corresponding letter in the Hebrew domain. Note that as the latter lacks vowels, a strict 1-to-1 alignment is not possible. Figure 2 illustrates the situation once the ***adhoc*** function has been allowed to create new lexical terms to cluster each pairing of letters under an analogically-specific category.

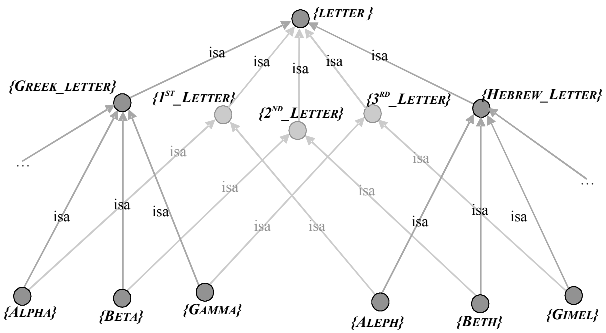

Figure 2: *WordNet supplemented with new adhoc categories like Greek-letter, Hebrew-letter and 1^{st}-letter, created as by-products of analogical retrieval.*

5.1. Evaluation

We first consider the effectiveness of ad-hoc category construction on the precision and recall of analogical retrieval in the WordNet deities domain. Table 1 presents the results of an experiment in which analogical variants are sought for the members of five different families of deity.

Table 1: *Cross-domain variants are sought for each member of five deity pantheons*

Ad-hoc categories	*Greek*	*Roman*	*Hindu*	*Norse*	*Celtic*
supreme-deity	Zeus	Jove	Varuna	Odin	N/A
wisdom-deity	Athena	Minerva	Ganesh	*n/a*	Brigit
beauty-deity, love-deity	Aphrodite	Venus	Kama	Freyja	Arianrhod
sea-deity*	Poseidon	Neptune	*n/a*	*n/a*	Lir
fertility-deity	Dionysus	Ops	*n/a*	Freyr	Brigit
Queen-deity	Hera	Juno	Aditi	Hela	Ana
war-deity*	Ares	Mars	Skanda	Tyr	Morrigan
hearth-deity	Hestia	Vesta	Agni	*n/a*	Brigit
Moon-deity	Artemis	Diana	Aditi	*n/a*	*n/a*
sun-deity*	Apollo	Apollo	Rahu	*n/a*	Lug

* *WordNet already contains the lexical concepts Sea-God, War-God and Sun-God*

This experiment thus involves 20 different mapping tasks (i.e., Greek to Roman deities, Hindu to Norse deities, Celtic to Greek deities, etc.). The average precision of analogical retrieval across all tasks is 93%, while the average recall is 61%.

For the letter mapping experiment, an analogous Hebrew letter was retrieved for each Greek letter, and vice versa. The ad-hoc categories created for each retrieval are of the form 1st-letter, 2nd-letter, and so on, and serve to pinpoint a precise analogue whenever one is available (that is, each ad-hoc category has an extension containing precisely two members). The average precision for the letter experiment is thus 100%. Since the Greek alphabet has more letters than the Hebrew alphabet, recall is 100% for the Hebrew to Greek task, but only 96% for the Greek to Hebrew task (since the latter has one less letter than the former).

5.2. Explicit Category Creation in WordNet

Though we have described the process of ad-hoc category creation as an implicit by-product of analogical reasoning, our formulations of ***alt*** and ***adhoc*** nonetheless allow us

to exploit analogy as a deliberate mechanism of explicit category and term creation. For every lexical concept A in WordNet, we need simply consider those alternate categorizations (derivable via ***alt***) that are also generated by at least one other concept:

$$\textbf{\textit{adhoc}}(\text{A:???}) = \{\text{X-Y} \mid \text{X-Y} \in \textbf{\textit{alt}}(\text{A}) \wedge \exists \text{B A} \neq \text{B} \wedge \text{X-Y} \in \textbf{\textit{alt}}(\text{B})\}$$

In effect, we are generating alternate categorizations of a given concept that have the analogical potential to relate that concept to at least one other in the ontology. That is, we interest ourselves here only with those alternate categorizations that possess an extension of two or more members, and which might thus make non-trivial additions to the ontological lexicon to serve a genuine organizational purpose. For example, the alternate categorization Greek-Wine constitutes a trivial addition to WordNet, since it serves to index a sole category member, Retsina. In contrast, the categorization Italian-Wine serves to index at least three different members (sweet vermouth, soave and Chianti). By this measure, Italian-Wine serves a useful indexing and clustering role in the ontology and should be retained, while Greek-Wine serves no clustering role and should be discarded[6]. In the food domain alone, WordNet provides definitions for over 200 different terms whose gloss mentions a proper-named country like “Italy”, “Greece” or “Mexico”, so we should expect that the alternate categorizer (as formulated via ***alt***) will pick out these national ties as features to be reified.

Applying the above formulation of ***adhoc*** to the 70,000+ noun concepts in WordNet, we obtain 8564 new and non-trivial compound categories. In total, these 8564 compounds differentiate 2737 different head concepts, suggesting that each head is differentiated in three different ways on average. Overall, the most differentiating modifier is “Mexico”, which serves to differentiate 34 different heads; for example,

[6] While Greek-Wine serves no useful clustering role, inasmuch as it serves to index just one concept, it might be seen as a useful addition to the ontology for reasons of symmetry. In an ontology that contains nodes like French-Wine, Italian-Wine and German-Wine, the addition of Greek-Wine, if only to index a single instance, would enhance the systematicity of the ontology. The criteria considered in this section should thus be viewed as heuristics rather than hard constraints.

Mexico-Dish serves to group together Taco, Burrito and Refried-beans. The most differentiated head is "herb", which is differentiated into 134 sub-categories such as Prickly-Herb, Perennial-Herb, European-Herb, etc.

To consider just a few other domains: sports are differentiated into team sports, net sports, court sports, racket sports and ball sports (surprisingly, but not meaninglessly, Bingo becomes categorized as a Ball-Game); constellations are divided into northern and southern variations; food dishes are differentiated according to their nationalities and their ingredients, e.g., into cheese dishes, meat dishes, chicken dishes, rice dishes, and so on. As noted earlier, letters are differentiated both by culture, giving Greek letters and Hebrew letters, and by relative position, so that "Alpha" is both a 1st_letter and a Greek_letter, while "Aleph" becomes both a 1st_letter and a Hebrew_letter. Likewise, Deity is further differentiated into War_deity, Love_deity, Wine_deity, Sea_deity, Thunder_deity, Fertility_deity, and so on.

One can legitimately ask whether such terms are truly creative, for it seems that we comprehend linguistic creativity here in its broadest sense, that used by Chomsky (1957) to describe the potential of human language to generate (i.e., create) an unlimited number of valid word combinations. It would seem that by our reckoning, then, that any novel combination of words that is syntactically and semantically valid should be considered creative. This criticism would certainly be apropos if the compounds under consideration were either entirely lexical or entirely conceptual. However, these compounds are both lexical *and* conceptual and are created relative to a lexical ontology in which they serve a useful organizational role. Compounds like "Strong-Drink" or "Love-deity" may seem mundane as linguistic artefacts in the context of general language usage, but from the context of a lexical ontology, they represent an insightful partitioning of a given conceptual space. Creativity requires

clarity of perception, and the value of this insight can be seen most forcefully in the kinds of analogies that these new categories allow one to construct. For instance, the category Strong-Drink creates a cluster of diverse (but appropriate) bedfellows from espresso (strong coffee) to concentrated orange juice (strong juice) to whiskey (strong liquor). In turn, this cluster provides a firm lexico-conceptual basis for analogies of the kind whiskey:liquor::espresso:coffee.

6. Conclusions

With this paper we have attempted to provide a common formalization – in terms of lexical composition and decomposition operators – for two different perspectives on the production of new lexical terms and categories. These perspectives are both analogically-motivated, and concern the explicit and implicit use of analogy in the implicit and explicit creation of new terms and categories. As such, it should be clear that we assign to analogy a central role in the mechanism of linguistic creativity.

Both perspectives create compound terms of the same form – simple modifier-head constructions – using lexical precedents to ensure that each term possesses both a linguistic and conceptual validity. However, while the outputs of both processes may look similar, the processes themselves are quite different, and cover different parts of the lexico-conceptual space. The implicit analogy approach, for instance, is only capable of generating compounds that can be reached via modulation from some existing support base in the lexicon. In contrast, the explicit analogy approach works directly with the feature-theoretic representation of concepts in the lexicon, and can generate compounds that, while meaningful, may have an empty support set. Each approach is thus complementary to the other, and both taken together yield a creative reach that is beyond either alone.

However, it is clear that a fusion of both perspectives does not provide full coverage of the lexico-conceptual space, even when this space is limited to that of simple modifier head constructions. Consider the terms “Gastropub” (a public house that serves restaurant-quality food), “metrosexual” (a heterosexual male with female grooming and fashion habits[7]) and “retrosexual” (a back-formation from “metrosexual” that describes the prototypical heterosexual male against which metrosexuality is defined as a reaction). These terms each combine a bound morpheme with a free morpheme, and while their structure is easy to analyse, it is extremely difficult to hypothesize an effective generation mechanism that does not simply combine every bound morpheme with every free morpheme in the lexicon. These compounds cannot be predicted either on the basis of existing compounds (via modulation and/or phonetic similarity[8]) or on the basis of conceptual features alone. Rather, since they are created for use in a particular communicative context, it is this context that provides the missing features that would make possible both the prediction of “Gastropub” and “Metrosexual” as valid words, possible as well as a broad ontological categorization of their meanings (e.g., that gastropub is a kind of public house, or that metrosexual is a type of heterosexual male)..

Since portmanteau words like “Gastropub” and “Metrosexual” and “metrosexual” comprise one of the most interesting varieties of modern lexical innovation (e.g., see Veale and O’Donoghue, 2001), it would be a shame if this were all that one could conclude. Most likely, there exists a middle ground in which these terms might be, if not predictable from lexical structure, then constrained by lexical structure to the extent that the addition of automatic corpus analysis (using the WWW, say) might allow a

[7] Intriguingly, the term “metrosexual” was first coined by the British journalist Mark Simpson in 1994 (see Simpson, 1994), but lay dormant for the rest of the decade. The term underwent a resurgence in popular culture when used in a New York Times article in 2003.

[8] It may be that “Gastropub” obtains some minor support from its phonetic neighbour “Gastropod” *after* it has been created using other means.

computational system the ability to harvest such novel terms and categories as they arise in a cultural setting. As such, the analysis framework described here should provide an adequate basis for interpreting novel portmanteau words if such words could be *harvested* automatically in lieu of being *predicted* automatically. One harvesting source that we are currently investigating is Wikipedia[9], an on-line open-source encyclopaedia that is constantly updated and modified by a veritable army of users. The popularity of Wikipedia makes it an ideal source from which to harvest new words as they gain prominence in the language, long before these words earn their place in conventional print dictionaries. For instance, Wikipedia offers a detailed entry for each of “Gastropub”, “Metrosexual” and “Retrosexual”, provides links between related terms, contains sufficient context to allow an automated system to construct an interpretation (e.g., the “Gastropub” entry mentions both public houses and gastronomy), and in some cases, provides a pertinent analogy to explain the term (e.g., Wikipedia helpfully points out that the Gastropub, as commonly conceived, is the English equivalent of the French brasserie).

Exploring term creation in the context of resources such as WordNet and Wikipedia, which blur the traditional distinction between dictionary and encyclopaedia, constitutes an on-going research .programme that is predicated on the belief that term creation is a scaleable phenomenon through which one can explore creativity in general. That is, the issues in term creation run the gamut from phonological to conceptual, involving terms that range from the mundane to the humorous to the wildly creative. It is our hope that an understanding of the processes that underlie term creation may thus lead to a deeper understanding of creativity overall, one that can be ultimately be exploited to build computational systems that exhibit genuine linguistic inventiveness .

[9] http://www.wikipedia.org

REFERENCES

Attardo, S, Hempelmann, C. F. and Di Maio, S. (2002). Script oppositions and logical mechanisms: Modeling incongruities and their resolutions. *Humor: International Journal of Humor Research,* 15-1, 3-46.

Baron, J. (1977). What we might know about orthographic rules. In S. Dornic (Ed.), *Attention and performance VI*. Hillsdale, NJ: Erlbaum.

Barsalou, L. (1983). Ad-Hoc Categories. *Memory and Cognition*, 11(3), 211-227.

Carpuat, M., Ngai, G. Fung, P., and Church, K.W. (2002). Creating a Bilingual Ontology: A Corpus-Based Approach for Aligning WordNet and HowNet. In the *proceedings of GWC 2002, the 1st Global WordNet conference*, Mysore, India.

Chomsky, N. (1957). *Syntactic Structures*. The Hague: Mouton.

de Bono, E. (1994). Parallel Thinking. Viking Press: London.

Dong, Z. (1998). Knowledge Description: What, How and Who? *The Proceedings of the International Symposium on Electronic Dictionaries*, Tokyo, Japan.

Freud, S. (1905) *Der Witz und seine Beziehung zum Unbewußten*. Leipzig, Vienna: Dueticke. (Reprinted Frankfurt am Main: Fischer. English edition,1976. Jokes and their Relation to the Unconscious. Harmondsworth: Penguin.)

Hayes, J. and Veale, T. (2005). Creative discovery in the lexical validation gap. *Journal of Computer speech and Language* (in press).

Hutton, J. (1982). *Aristotle's Poetics*. New York, NY: Norton.

Lenat, D. and Guha, R. V. (1991). *Building Large Knowledge-Based Systems*. Reading, Massachusetts: Addison Wesley.

Pustejovsky, J. (1991). The generative lexicon. *Computational Linguistics,* 17(4).

Miller, G. (1995). WordNet: A Lexical Database for English. *Communications of the ACM*, 38(11).

Rumelhart, D. E. and Abrahamson, A. A. (1973). A model for analogical reasoning. *Cognitive Psychology*, 5, pp 1 – 28.

Simpson, M. (1994). "Here come the mirror men". *The Independent* Newspaper (London), November 15 edition, p. 22.

Trask, R. L. (1996). *Historical Linguistics*. London: Edward Arnold.

Veale, T. and O'Donoghue, D. (2001). Computation and Blending. *Cognitive Linguistics* 11(3/4), 253-281.

Veale, T. (2003a). Dynamic Type Creation in Metaphor Interpretation and Analogical Reasoning. *In the proceedings of the International Conference on Conceptual Structure, Conceptual structures for Knowledge Creation and Communication*, LNAI 2746. Berlin: Springer Verlag.

Veale, T. (2003b). The Analogical Thesaurus: An Emerging Application at the Juncture of Lexical Metaphor and Information Retrieval. *In the Proceedings of IAAI'03, the 2003 Innovative Applications of Artificial Intelligence conference*. Menlo Park, California: AAAI Press.

Veale, T. (2004). Creative Information Retrieval*. In the proceedings of CICLing'04. Lecture Notes in Computer Science*. Berlin: Springer Verlag.

Veale, T., Seco, N. and Hayes, J. (2004). Creative Discovery in Lexical Ontologies. *In the proceedings of COLING'2004, the 20th International Conference on Computational Linguistics*. Geneva, Switzerland. San Mateo, California: Morgan Kaufmann.

Veale, T. (2005). A Dynamic Type Hierarchy for Linguistic Creativity. *Journal of Knowledge-Based Systems* (in press).

Way, E. C. (1991). *Knowledge Representation and Metaphor*. Studies in Cognitive systems. Amsterdam: Kluwer Academic Publishers.

Re-Representation and Creative Analogy:

A Lexico-Semantic Perspective

Tony Veale,

School of Computer Science and Informatics,

University College Dublin,

Ireland.

Tony.Veale@UCD.ie

Abstract

Analogy is a powerful boundary-transcending process that exploits a conceptual system's ability to perform controlled generalization in one domain and re-specialization into another. The result of this semantic leap is the transference of meaning from one concept to another from which metaphor derives its name (literally: *to carry over*). Such generalization and re-specialization can be achieved using a variety of re-representation techniques, most notably abstraction via a taxonomic backbone, or selective projection via structure-mapping over propositional content. In this paper we explore both the extent to which a bilingual lexical ontology for English and Chinese, called HowNet, can support each technique, and the extent to which both are, ultimately, variations of the same process of creative re-representation.

Keywords: Analogy, Abstraction, Structure-Mapping, Ontologies, Lexical Semantics

1 Introduction

Given a recalcitrant category that unduly limits one's actions, the creative individual seeks a new category within which to operate. Indeed, this ability to reconceptualize an object or idea from a different perspective, or from within the viewpoint of a different category, is conventionally considered central to human creativity. Analogy is just one manifestation of creative cognition, but one that clearly illustrates this ability of creative thinkers to transcend conceptual boundaries and perform a semantic leap from one category structure to another (e.g., see Veale, 2003,2004b). Reconceptualization can be dramatic, as when a scientist is forced to accept a paradigm shift from one theory to another (see Kuhn, 1962), or mundane, as when someone uses a credit-card to open a door, a screw-driver to open a can of paint, or a chair to wedge a door shut. Humour also employs reconceptualization as a resolution mechanism: consider how many jokes stretch or transcend the definitional boundaries of conventional categories[1], or employ a punch-line that forces a listener to recategorize his or her interpretation of the preceding narrative (e.g., see Ritchie, 1999; Attardo et al., 2002; Veale, 2004a).

Reconceptualization is certainly a good high-level *story* of what occurs in creative situations, but as computationalists, we require a more specific account. In particular, we require an algorithmic insight into what it means to transcend category boundaries, and this in turn requires some minimal commitment to some form of conceptual

[1] For instance, many jokes play with the boundaries of taboo categories to categorize non-taboo events – like visiting the doctor – in terms of taboo events such as a sexual infidelity (see Attardo *et al*., 2002). Alternately, some jokes re-imagine socially mediated categories, like those that constitute our value systems, in subjective terms. Consider the following remark from a famously talented, and famously dissolute footballer: "I spent most of my money on alcohol, women and gambling, and the rest I wasted."

representation. In fact, computationalists of an empirical bent often prefer to work from the latter to the former: given a particular resource with a specific conceptual representation, empiricists attempt to frame the problem in terms of this representation. Perhaps unsurprisingly, this is precisely what we shall attempt to do in this paper. The resource we focus on here is HowNet, a bilingual lexical ontology for Chinese and English (see Dong, 1988; Wong, 2004). Earlier experiments (e.g., see Veale, 2004b, 2005) suggest that HowNet is well suited to the demands of analogy reasoning, and thus some forms of creative reasoning, since HowNet combines a taxonomic backbone with an explicit, if somewhat sparse and under-specified, propositional semantics.

Now, theories of analogy and metaphor are typically based either on structure-mapping (e.g., see Falkenhainer *et al.*1989; Veale and Keane, 1997) or on abstraction (e.g., see Hutton, 1982; Fass, 1988; Way, 1991; Veale, 2003). While the former is most associated with analogy, the latter has been a near-constant in the computational treatment of metaphor. Structure-mapping assumes that the causal behaviour of a concept is expressed in an explicit, graph-theoretic form so that unifying sub-graph isomorphisms can be found between different propositional representations. In contrast, abstraction theories assume that analogous concepts, even if far removed in ontological terms, will nonetheless share a common hypernym that will capture their causal similarity. Thus, we should expect an analogous pairing like *cancer* and *assassin* to have very different immediate hypernyms but to ultimately share a behavioural abstraction like *kill-agent* (e.g., see Veale, 2003).

With a well known lexical ontology like WordNet (see Miller, 1995), both structure-mapping and abstraction-based approaches are problematic. The idea that a

one-size-fits-all representation like WordNet will actually provide a hypernym like *kill-agent* seems convenient almost to the point of incredulity. As much as we want our ontologies to anticipate future analogies with these pro-active categorizations, most off-the-shelf ontologies simply do not possess such convenient terms (see Wong, 2004). Similarly, WordNet lacks the propositional content that is the necessary grist for a structure-mapping approach. The semantic content that would ideally fill this role is not explicit, but implicitly resides in the unstructured textual glosses that annotate each lexical concept.

In this paper we explore the extent to which another lexical ontology, the aforementioned Chinese/English HowNet system (see Dong, 1988; Carpuat *et al*. 2002; Wong, 2004) supports the kind of reconceptualization that is required in the generation and interpretation of creative analogies. The WordNet-like taxonomic backbone, in combination with its own unique propositional semantics, allows us to evaluate the extent to which both structure-mapping and abstraction theories of analogy can be supported by the same lexical ontology.

We begin by briefly summarizing past approaches to the computational treatment of metaphor and analysis in section 2, before comparing the pros and cons of WordNet and HowNet in section 3. In section 4 we describe a form of reconceptualization that relies on conceptual abstraction; however, we do not propose a model of simple taxonomic abstraction, but one of *relational* abstraction, since only the latter allows us to generalize over the functional and behavioural meaning of a concept. To extend the reach of relational abstraction to representations that would not otherwise support this technique, we also present here a form of representational transformation called *structural inversion*. This is, in essence, a form of figure-ground reversal in which alternative representations for an under-specified concept can be sought by turning to elements in the conceptual background. In section 5 we then describe a form of reconceptualization based on structural rarefaction; this in turn supports a structure-mapping approach to analogy using HowNet representations. In section 6, both of these approaches to reconceptualization are subjected to a comparative evaluation across the entirety of HowNet. We conclude by arguing, on the basis of this evaluation, that these approaches are ultimately complementary, inasmuch as a synthesis of both produces better performance than does either approach in isolation.

2 Past Work

That analogy and metaphor operate across multiple levels of conceptual abstraction has been well known since classical times. Aristotle first provided a compelling taxonomic account of both in his Poetics (see Hutton, 1982 for a translation), and computationalists have been fascinated by this perspective ever since. While the core idea has survived relatively unchanged, one must discriminate theories that apparently presume a static type-hierarchy to be sufficient for all abstraction purposes (e.g., Fass, 1998), from theories that posit the need for a dynamic type hierarchy (e.g., Way, 1991; Veale, 2003). One must also differentiate theories that have actually been implemented (e.g., Fass, 1988; Veale, 2003,2004) from those that are either notional or that seem to court computational intractability (e.g., Hutton, 1982; Way, 1991). Perhaps most meaningfully, one must differentiate theories and implementations that assume

hand-crafted, purpose-built ontologies (e.g., Fass, 1988) from those that exploit an existing large-scale resource like WordNet (e.g., Veale, 2003,2004). The latter approach side-steps any possible charge of hand-crafting by working only with third-party resources, but at the cost of living with their perceived flaws and inadequacies.

Structure-Mapping theory is founded on the premise that the most satisfying analogies are those that operate at the causal level of representation, since causality allows an analogy to offer a deep explanation for a poorly understood phenomenon (e.g., see Falkenhainer *et al*. 1989). Thus, *the atom as miniature solar-system* is a satisfying analogy because both source and target are causally structured around the notion of rotation. Furthermore, when comparing agents or artefacts (e.g., see Veale and Keane, 1997), this causality can be captured by considering the functional or behavioural commonality between target and source: a footballer can be meaningfully described as a gladiator or a warrior since each exhibits competitive behaviour, and a scalpel can be compared to a sabre, a sword or a cleaver since each has a cutting behaviour.

By employing a single lexical resource, HowNet, to implement both the relational abstraction *and* the structure-mapping theories of analogy, we have as a secondary goal a demonstration that both perspectives are not fundamentally opposed. Structure-mapping can be seen as a form of structural-abstraction, where one abstracts out the causal backbone of a concept, while taxonomic abstraction, if performed upon the relations implied by a concept rather than the concept itself, can also be seen as a highly selective form of structure-mapping. Ultimately, both kinds of approach attempt to capture the functional or behavioural commonality between a pair of source and target concepts: a surgeon can be meaningfully described as a repairman since both occupations have the function of restoring an object to an earlier and better state; the distinction, which is glossed over both by abstraction and structure-mapping approaches, is that a surgeon restores by healing, while a repairman restores by mending.

3 Comparing WordNet and HowNet

HowNet and WordNet each reflect a different view of semantic organization. WordNet is *differential* in nature: rather than attempting to express the meaning of a word explicitly, WordNet instead differentiates words with different meanings by placing them in different synonym sets, and further differentiates these synsets from one another by assigning them to different positions in its taxonomy. In contrast, HowNet is *constructive* in nature. It does not provide a human-oriented textual gloss for each lexical concept, but instead combines sememes from a less discriminating taxonomy to compose a semantic representation of meaning for each word sense.

For example, the lexical concept *surgeon|*医生 is given the following semantic definition in HowNet:

surgeon|医生≡{human|人:HostOf={Occupation|职位},
domain={medical|医}},
{doctor|医治:agent={~}}}

which can be glossed thus: "a surgeon is a human with an occupation in the medical domain who acts as the agent of a doctoring activity." The *{~}* construct serves as a

self-reference, to mark the location of the concept being defined in the given semantic structure. The oblique reference offered by the tilde serves to make the definition more generic, so that many different concepts can conceivably employ the same definition. Thus, HowNet uses the above definition not only for surgeon, but for medical workers in general, from orderlies to nurses to internists and neurologists.

Perhaps because HowNet relies less on hierarchical differentiation, it has a considerably less developed middle ontology than WordNet. For instance, most kinds of person in HowNet[2], from mathematicians to hobos, are placed directly under the hypernym *human|*人, eschewing the intermediate concepts like *{professional}, {specialist}* and *{worker}* that give substance to WordNet's middle ontology. We note that HowNet does indeed define these concepts – but unlike WordNet, it does so at the leaf level where they add nothing to the internal structure of the taxonomy.

3.1 Analogical Signatures and HowNet

Nonetheless, the skeletal nature of HowNet semantic definitions, combined with the wide-spread use of *{~}* as a generic reference, suggests how HowNet might support an efficient approach to analogical reasoning. By indexing each concept on a reduced form of its semantic definition – an *analogical signature* – analogies will correspond to collisions between concepts with different definitions but with identical signatures. Such an approach can be efficiently implemented using simple string hashing of signatures, to

[2] We note in passing that the Chinese origins of HowNet explains some additional, cultural distinctions between Princeton WordNet and HowNet. For instance, WordNet defines dogs as a kind of canine; HowNet defines dogs as kinds of livestock.

detect analogical collisions between kitchens and factories, generals and admirals, ballet dancers and acrobats, or cruise missiles and arrows. The devil here is in the *lack* of detail: because HowNet's definitions are frequently imprecise and fail to fully specify a concept, they allows others – potential analogues – to occupy the same reduced semantic space. The further we exacerbate this deficiency, indexing each definition on an increasingly diluted version of itself, the more distant and creative will be the analogies that are generated. For example, excluding the hypernym of a definition, or its domain markings, facilitates analogies between people and non-people, such as pests and persecutors, or hackers and viruses.

To implement both the abstraction and structure-mapping theories of analogy, we will explore the effectiveness of two kinds of analogical signatures in this paper: relational signatures derived, via abstraction, from the predicate and case-role of a proposition, and structured, template-like signatures based on generalized propositional content in which place-holder variables may be added.

4 Re-Representation via Abstraction Signatures

Given the general impoverishment of HowNet's middle ontology (at least compared with that of WordNet), abstraction-based signatures should not be based directly on taxonomic organization. Rather, by instead deriving analogical signatures from the relational structure of a concept's semantic definition, we can better capture the functional and behavioral nature of the concepts concerned. We can do this by focusing on how each concept is situated with respect to its relational context, which is to say, by

targeting the explicitly self-referential *{~}* in each definition. For instance, consider the following semantic definition of repairman in HowNet:

repairman|修理工 ≡{human|人:
HostOf={Occupation|职位},
{repair|修理:agent={~}}}

Noting the relational position of *{~}* here, we can infer that a repairman is the agent of a repairing activity. Expressing this as a new taxonomic type, we can reify the combination of activity and role to create a new taxonomic term *repair-agent*, of which repairman will be an instance. From an analogical perspective, *repair-agent* thus serves as a good relational signature for *repairman|修理*.

Further noting that the HowNet taxonomy defines the predicate *repair|修理* as a specialization of the reinstatement predicate *resume*|恢复, we can further establish *repair-agent* as a specialization of *resume-agent*[3]. This double layer of abstraction effectively establishes a new, parallel taxonomy that organizes lexical-concepts according to their analogical potential, rather than their formal taxonomic properties. For instance, as shown in Figure 1, *resume-agent* will encompass not only *repair-agent*, but *doctor-agent*, since HowNet also defines the predicate *doctor|医治* as a specialization of *resume*|恢复.

[3] HowNet uses the predicate *resume* in the sense of *restore*, that is, "to resume an earlier, better state".

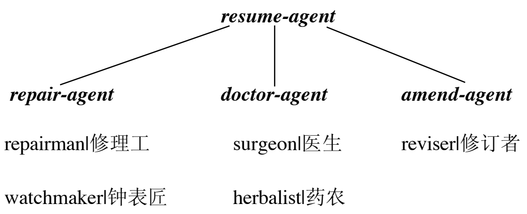

Figure 1: Portion of a new three-level abstraction hierarchy derived from HowNet's relational structures.

Relational signatures like *repair-agent* and *doctor-agent* are, in essence, new lexical concepts that allow particular problems of an analogical or metaphoric nature to be solved creatively. Hierarchies like that of Figure 1 thus reflect the general philosophy of creativity espoused in McCarthy (1999), which stipulates that a solution is only truly "creative" when it recruits or creates concepts that were not directly mentioned in the original problem specification. McCarthy's viewpoint is interesting because it forces us to evaluate creativity not just on the utility of the end-product, which might equally be produced by the most banal of exhaustive searches, but on the selective means through which this end was achieved.

In general, relational signatures are generated as follows: given a semantic fragment *F:role={~}* in a HowNet definition of a concept C, we create the signatures *F-role* and *F'-role*, where *F'* is the immediate HowNet hypernym of *F*, which in turn is the immediate hypernym of *C*. The role in question might be *agent, patient, instrument,* or any other role supported by HowNet, such as *target, content*, etc.

Each concept is thus assigned two different relational signatures: a direct signature

(*F-role*) based on the specific relational structure of the concept, and another more abstract signature (*F'-role*) that is generalized, via taxonomic abstraction, from this direct signature. These signatures effectively form an alternate taxonomy by which the lexical concepts in HowNet can be organized for analogical purposes. Figure 2 illustrates a partial hierarchy derived from HowNet definitions of form-altering tools:

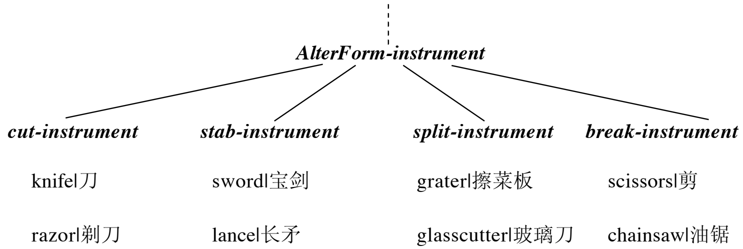

Figure 2: a derived taxonomy of relational signatures that facilitates analogy between instruments that "alter the form" of other objects.

This additional layer of abstraction is necessary to facilitate creative analogy between semantically distant concepts. Nonetheless, we note that since HowNet's designers have already exercised a certain degree of metaphoric license, even concepts with the same direct signature can exhibit a surprising degree of semantic variety.

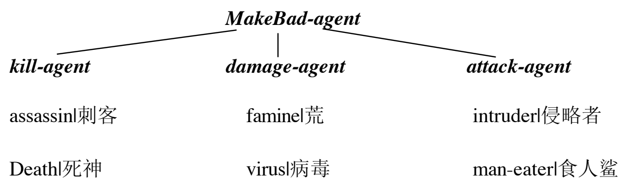

Figure 3: semantic diversity among concepts with the same relational signatures.

This diversity, as illustrated by Figure 3, means that the analogy "Death is an assassin" can be generated in a single generalization step, while the analogy "Death is a man-eater" can be generated with just two generalization steps.

4.2 Reconceptualization via Structural Inversion

Since the partial taxonomies of Figures 1, 2 and 3 do not exist in HowNet, but are derived from HowNet representations, it seems quite meaningful to refer to these taxonomies as reconceptualizations of the original HowNet taxonomy. However, the technique of relational abstraction is seriously limited, in a way we shall explicitly quantify later, by its ability to apply only to those definitions that are self-referential. If no relational signature can be generated for a given concept, as is the case when a definition is not structured around the use of {~}, then no analogies can be retrieved for that concept. For example, consider the HowNet definition of "*bicycle*":

bicycle|单车 ≡ {LandVehicle|车: modifier={manual|非自动}}

Clearly, no analogical signature can be derived from this overly under-specified definition. Nevertheless, consider another HowNet entry that refers to this *bicycle|单车* definition:

cyclist|车手 ≡

{human|人:

{drive|驾驭:

agent ={~},

patient=<u>{LandVehicle|车: modifier={manual|非自动}}</u>}}

The concept *cyclist|车手* is clearly much better situated with respect to analogical reasoning, giving rise to the signatures *drive-agent* and its abstraction *CauseToMoveInManner-agent* that are shared by pilot, chauffeur, astronaut and trucker amongst others. Since bicycle occupies the conceptual background of this definition, a figure-ground reversal is needed to bring it into the foreground as the focus of the definition. That is, we can structurally invert this definition to yield an alternate conceptualisation of *bicycle|单车*, by simply replacing the {~} marker with the foreground concept *cyclist|车手* for which it stands in the definition, thus backgrounding this concept, and replacing the sub-definition of *bicycle|单车* with {~}, thus foregrounding this concept. This figure-ground reversal is graphically illustrated in Figure 4:

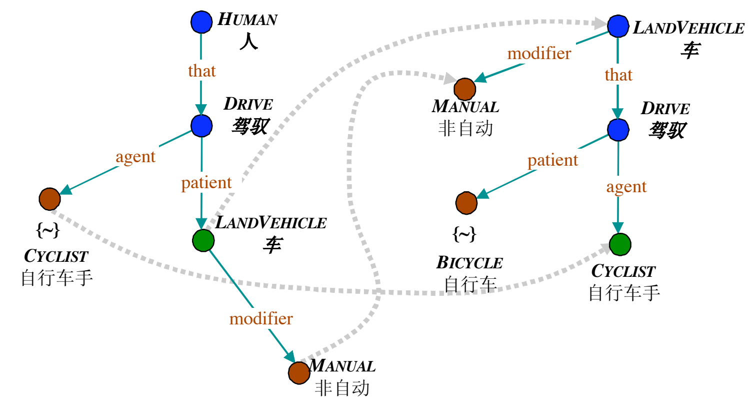

Figure 4: The HowNet definition of *cyclist|车手* is structurally inverted to yield a richer representation of *bicycle|单车* than that which is offered by HowNet itself.

Further replacing the taxonomic head of the definition (human|人) yields this new structure:

bicycle|单车 ≡

{LandVehicle|车:

{drive|驾驭:

agent ={*cyclist|车手*},

patient={~}}}

In this reconceptualization, a bicycle is a vehicle that is driven by a cyclist. While this is not a particularly tight definition, it is precisely this lack of formal rigidity that serves to enable creative thinking. After reconceptualization, the concept *bicycle|单车* is thus assigned the analogical signature *drive-patient*, facilitating an analogical mapping to boats, airplanes, trucks and even elevators (since HowNet construes each as the patients of a driving activity).

5 Re-Representation via Structural Signatures

The structure-mapping approach also strives for abstraction, not through the selective creation of new taxonyms but through a form of structural *rarefaction*. Recall that structure-mapping theory places particular emphasis on the causal backbone of a concept's propositional content, which is usually projected unchanged from one domain to another (see Falkenhainer *et al*. 1989). Based on this isomorphic alignment of relational structures, the entities contained in each structure are typically placed into a 1-to-1 correspondence with one another. The attributive modifiers of these entities play a more peripheral role in structure-mapping, but in approaches like *Sapper* (Veale and

Keane, 1997) they often serve as a literal grounding for an analogy. Figure 5 depicts an example of the structure-mapping process applied to HowNet representations.

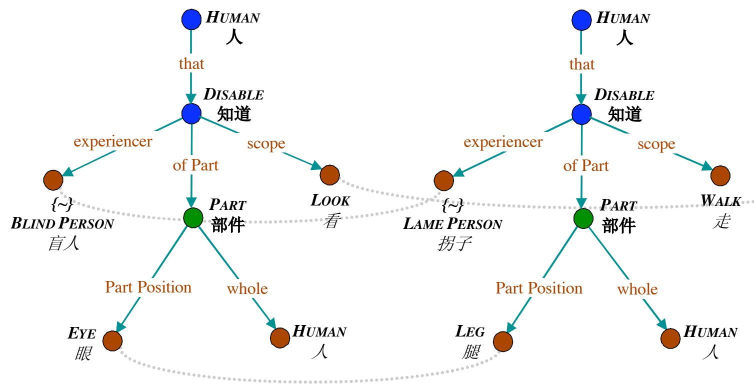

Figure 5: Structure-Mapping applied to the HowNet representations of *blind person|盲* and *lame person|拐子*, identifying the 1-to-1 entity mappings *look*:*walk* and *eye*:*leg*.

In many cases, the semantic definitions provided by HowNet are already so skeletal and under-specified that we may assume that any structure-mapping signature will preserve the general form or shape of the proposition from which it is derived, one signature per proposition. As an example, consider the HowNet definition of *blind person|盲人*:

{human|人: {disable|知道:

OfPart={part|部件:PartPosition={eye|眼}, whole= {human|人}}

experiencer={~},

scope={look|看}}}

In other words, a blind person has "a disability of the eye that affects one's ability to look". One finds precisely the same propositional structure in the HowNet definition of *lame person|拐子*, except that *eye|眼* is replaced with l*eg|腿* and *look|看*is replaced with *walk|走*. The goal of a structure-mapping approach is to capture precisely this semantic isomorphism while simultaneously identifying entity-level differences like *eye:leg* and *look:walk* as cross-domain counterparts. We thus need to generalize from each proposition in every definition a structural signature that, by virtue of being identical to another, signals a structural equivalence between the underlying definitions. For instance, the shared signature for *blind person|盲人* and *lame person|拐子*looks like:

{?:{ill|病态:OfPart={?},experiencer={~},scope={?}}}

Generalized structural signatures of this form are generated via a 7-step process:

1. Split each definition into multiple propositions, and generate a separate signature for each.
2. If a proposition describes a noun concept, replace its taxonomic head with a *?* marker (E.g., human|人 → ?). In contrast, if a proposition describes a verb concept, replace its taxonomic head with its most specific hypernym (e.g., *repair* → *resume*).
3. Replace the conceptual arguments bound to each case-role of a predicate with the variable marker *{?}*. These markers will indicate positions in the signature where 1-to-1 correspondences between source and target structures can be made.
4. When a propositional sub-structure corresponds to the definition of another concept, replace the entire sub-structure with a *{?}* variable marker as in 3 above.

5. Replace predicates by their immediate hypernyms in the HowNet taxonomy. Thus, both *repair|修理* in the definition of *repairman|修理*, and *doctor|医治* in the definition of surgeon|医生, should be replaced by the hypernym *resume*|恢复 when generating their respective signatures.
6. Remove any explicit domain tag in a proposition from the corresponding signature (e.g., the assignment *domain={medical|医}* in the definition of *surgeon|医生*). This is necessary since analogy is meant to transcend domain boundaries.
7. Generalize the value of any purely attributive relation, like *modifier, manner, restrictive, host* or *content*, to its immediate hypernym, and ensure that step 3 above does not variablize the resulting value but allows it to remain present as a literal.

Not all of these steps need to applied to produce a valid signature. For instance, multiple signatures at different levels of detail can be generated for the same proposition by alternately applying or ignoring steps 4 and 5. Indeed, because a different signature is generated for each sub-proposition (except for empty propositions, as we shall discuss below), a given HowNet definition will often generate several structural signatures, so that overall, there may be more unique signatures than unique propositional structures.

Following these 7 steps then, the following structural signatures will be assigned to each of the concepts *surgeon|医生, repairman|修理, reviser|修订者, watchmaker|钟表匠* and *herbalist|药农*:

{?:HostOf={?}} and *{?: resume*|恢复*:agent={~}}}*

However, because the *HostOf* relation always occurs with the binding *Occupation|职位* in HowNet, it is effectively useless as an analogical index and the resulting signature is discarded. So in the example above, only the latter signature is retained.

More structural richness is exhibited by the lexical concepts *apostle|使徒* and *insider|局内人*, whose HowNet definitions are shown below.

apostle|使徒

≡ {human|人:

{believe|修理:

agent={~},

content={humanized|拟人},

domain={religion|宗教}}}

person who knows inside story|个中人

≡ {human|人:

{know|知道:

agent={~},

content={fact|事情:

modifier=covert|隐秘}}}

These are also assigned the same structural signature:

{?: {HaveKnowledge|有知:agent={~}, content={?}}

The sub-structure *{fact|事情:modifier=covert|隐秘}* is completely variablized within the signature of *person who knows inside story|个中人* since this corresponds to the HowNet definition of *secret|秘事* (see step 4). Analogically then, an apostle is a religious insider, one who knows the inside scoop on a given deity (denoted *humanized|拟人* in HowNet).

6 Comparative Evaluation

Consider first the composition of the HowNet version used in this research. It contains 95,407 unique lexical concepts (excluding synonyms) and 23,507 unique semantic definitions. Clearly then, these definitions are under-specified to the extent that many are shared by non-identical concepts (such as *cart|板车* and *bicycle|单车*, which HowNet simply defines as manual vehicles with the same under-specified definition). Furthermore, 90% of these definitions comprise a single proposition, while only 8% comprise two propositions and a mere 2% comprise three or more propositions.

We evaluate the abstraction and structure-mapping approaches using four criteria: *coverage* – the percentage of unique HowNet definitions from which a valid signature can be derived; *recall* – the percentage of unique definitions (not concepts) for which at least one analogical counterpart can be found; *parsimony*– the percentage of effective signatures that can actually be used to generate analogies (the most parsimonious approach is precise in generating only those signatures that are analogically useful); and *richness* – the complexity of the mappings captured by each analogy, as measured by the average number of entity correspondences per analogy.

6.1 Evaluating Relational Abstraction

6.1.1 Abstraction Coverage

Since relational signatures exploit occurrences of *{~}* for their generation, both the coverage and recall of the relational abstraction approach depend crucially on the wide-spread usage of this reflexive construct.

However, of the 23,507 unique definitions in HowNet, just 6430 employ this form of self-reference. The coverage offered by relational signatures is therefore just 27% of the available definitions. However, structural inversion enlarges the HowNet semantic space from 23,507 unique definitions to 24,514, with each of these additional 1007 definitions employing {~} self-reference. The coverage of analogical mapping with structural inversion is thus 31% (which represents a 15% improvement).

We note that while 31% is still rather low, the use of {~} is not uniform across HowNet's definitions. The most useful concepts from an analogical perspective, Person, Animal and Artefact, are more densely represented by self-referential definitions than the ontology as a whole, offering 65%, 68% and 42% coverage respectively.

6.1.2 Abstraction Recall

From those definitions containing a {~} self-reference, 1579 unique direct signatures are generated. In turn, another 838 abstracted relational signatures are generalized from

these using HowNet's taxonomic organization of verbs. In total, 2219 unique relational signatures are generated. This reveals that the sets of direct and abstracted signatures are not disjoint, and that in 8% of cases, the abstracted signature of one definition corresponds to the direct signature of another.

The overall recall rate is 30% (or 26% without structural inversion), which is to say, a relational signature enables the recall of at least one analogous definition for 30% of the unique definitions in HowNet. The most productive relational signature is *control_agent*, which serves to analogically co-index 210 unique HowNet definitions, among them the definitions of Boss, Manager, Manipulator, Bosun and Traffic-Cop.

6.1.3 Abstraction Parsimony/Precision

Overall, 1,315 of all 2219 relational signatures prove to be useful in co-indexing two or more definitions, while 904 relational signatures are associated with just a single definition. The parsimony of the abstraction approach is thus 59%, which is to say that 59% of the generated signatures are analogically useful, while 41% serve no analogical purpose and are ultimately rejected. This measure of parsimony is a useful index of predicate re-use in HowNet: a high parsimony score suggests that most definitions are defined using a communal set of predicates that systematically apply to more than one concept; a low parsimony score suggests that most definitions are defined on an ad-hoc basis. A parsimony score of 59% is moderate, suggesting strong systematicity but some ad-hoc tendencies in HowNet.

6.1.4 Abstraction Richness

Since the abstraction approach produces atomic, rather than structured signatures, it is capable of generating only one mapping per analogy, at the gross level of the source and target concepts themselves. For instance, while the abstraction approach can recognize that *blind person|盲人* and *lame person|拐子* are analogous by virtue of sharing the relational signature *disable-experiencer*, it cannot recursively determine entity mappings like *eye:leg* and *look:walk* in the way that structure-mapping can. The taxonomic approach thus has a uniform mapping richness of 1.

6.2 Evaluating Structure-Mapping

6.2.1 Structure-Mapping Coverage

A structure-mapping signature can be generated for every structured definition in HowNet. In principle then, the coverage of this approach is 100%. In practice, however, 10% of HowNet's semantic definitions contain no real structure beyond the specification of a hypernym or a domain tag. The maximum coverage of structure-mapping then, as limited to definitions with relational structure, is 90%.

6.2.2 Structure-Mapping Recall

HowNet's 21,761 unique structured definitions comprise 21,929 unique propositions. From these, 21,159 unique structural signatures are derived (many of which are generalizations of other signatures), serving to find analogues for 14,370 definitions. The recall of structure-mapping is thus 61%, while the most productive signature is:

{component.部分:whole={?}}

This signature serves to analogically co-index the 397 unique definitions for concepts that exhibit a part-whole distinction.

6.2.3 Structure-Mapping Parsimony/Precision

With 79% of all structural signatures serving to index just a single definition, the parsimony of the structure-mapping approach must be judged as a low 21%. However, the parsimony of the structure-mapping approach does not have the same critical import for HowNet's overall design as does the parsimony of the relational abstraction approach. Here we measure the reusability of structural forms, or patterns, rather than the predicates that semantically anchor these forms. Nonetheless, a higher parsimony score is desirable, and would reflect a higher degree of structural organization in HowNet.

6.2.4 Structure-Mapping Richness

Most analogies (64%) generated using the structure-mapping approach imply two entity mappings, 25% imply three entity mappings, and 11% imply four or more. The average mapping richness of a structure-mapped analogy is thus 2.48.

6.3 Analysis of Results

The results of this comparison, as summarized in Table 1 below, force us to draw some important conclusions about the utility of each approach to performing analogical reasoning in HowNet.

	Abstraction	Structure-Mapping	Combination
Coverage	.31	.90	.90
Recall	.30	.61	.72
Parsimony	.59	.21	.24
Richness	1.0	2.48	2.24

Table 1: Comparison of both approaches to analogy in HowNet

First, though the abstraction approach is capped by the limited use of self-reference among HowNet definitions, it demonstrates a recall rate that closely approaches this ceiling, managing to find analogies of non-trivial complexity for almost 1 in 3 HowNet definitions (or 1 in 4 without structural inversion). Because of its broader coverage, structure-mapping does considerably better, generating analogies for 3 in 5 definitions. A combination of both approaches ("combination" in Table 1) generates analogies for almost 3 in 4 definitions, which is most encouraging given the creative demands of analogy generation. This is especially so as we have considered here analogies between unique definitions, not unique words. The inherent ambiguity of natural language means that just one inter-definition analogy might be lexically realized in tens, perhaps even hundreds, of different ways.

7 Conclusions

Relational abstraction, structural inversion and structure-mapping are all forms or reconceptualization, since each derives new semantic structures from old. This paper has

explored these three different, but ultimately complementary, approaches to reconceptualization within the specific context of HowNet, a large-scale conceptual resource.

Though reconceptualization involves representational change, it is debatable whether reconceptualization as explored here possesses the radical power of conceptual change ascribed to the process of transformational creativity by Boden (1990, 1999). These approaches to reconceptualization do not dramatically reconfigure the conceptual space in which creative processing is to occur, but the conceptual space is nonetheless modified in an important way. For instance, the first approach, relational abstraction, allows a new taxonomic organization to be constructed from the relational *predicate:case-role:filler* structure of HowNet's propositional semantics. This new taxonomy is based not on conventional categories of being, but on the functional and behavioral nature of the concepts involved. As such, it provides multiple layers of relational abstraction that can facilitate creative analogy between semantically distant, but functionally similar, HowNet entries. This strategy can additionally be seen as a form of *meta*-reconceptualization, since it allows the constructive semantics of HowNet to be construed in the differential manner of WordNet. That is, through the process of relational abstraction, HowNet's conceptual space is altered such that HowNet's semantic structure becomes both constructive and differential. In this light at least, relational abstraction is a transformative process.

The second approach to reconceptualization is based not on abstraction, but on structural rarefaction, though philosophically, both mechanisms are similarly motivated. To the extent that a relational abstraction is based on the combination of a predicate and a case-role, it can be seen as a compressed and very regular form of structural signature. As such, we begin to appreciate that the abstraction and structure-mapping approaches to analogy are not that different after all. Both aim to reconceptualize a concept in a way that allows important semantic similarities to be highlighted, while unimportant dissimilarities are forced into the background.

The third mechanism of reconceptualization we have explored is structural inversion, which effectively allows a system to look outside a concept to obtain a new semantic perspective from the vantage point of other concepts. The strategy of structural inversion clearly complements that of relational abstraction, since the former provides additional propositional content for the latter to abstract over. In fact, structural inversion often provides multiple alternate perspectives on a concept, any of which might be used to generate an analogy or, more generally, to solve a problem. Consider the concept *software|软件*, which HowNet simply defines as a kind of *implement|软件*. Structural inversion allows software to be redefined, among other things, as anything that can compiled via a programming language, or anything that is damaged by a computer virus. This form of redefinition is clearly quite liberal, as not everything affected by a virus is generally deserving of the label "software". Yet, liberal categorization lies at the root of creative thinking: this redefinition forces us to consider web-pages, spreadsheets and even email messages as software, and indeed, under closer examination, all do fit the bill as "soft" wares.

The liberality of structural inversion seems well-suited to the robust treatment of categories whose membership criteria are arbitrary or highly subjective. Consider the concepts treasure, curio and oddity, each of which receive cursory treatment in HowNet's semantics. Structural inversion allows a system to reconceptualize the concept *treasure|珍宝* as anything that is stored in a jewellery box, sold in a jewellery shop, or hidden on a treasure

ship. Were someone to store something of subjective value in a jewellery box, such as love letters, photos, etc., a creative system based on structural inversion would certainly be able to recognize their value. Though the relative contribution of structural inversion to the "bottom line" of the HowNet evaluation is relatively slight, enabling the recall of the relational abstraction approach to jump from 26% to 31%, we believe it to be a promising technique that deserves further research in the context of other resources and creative tasks.

In closing, we note that the results of this work, in particular the perspective of relational abstraction, can be tangibly appreciated in the *Analogical Thesaurus*, an on-line index derived from HowNet that allows word-concepts to be retrieved using both analogy and metonymy. This index is available for use on-line at: *Afflatus.ucd.ie*.

References

Attardo, S, Hempelmann, C. F. and Di Maio, S. (2002). Script oppositions and logical mechanisms: Modeling incongruities and their resolutions. *Humor: International Journal of Humor Research* 15-1, 3-46.

Boden, M. (1990). *The Creative Mind: Myths and Mechanisms*. London: Weidenfeld and Nicholson.

Boden, M. (1999). Computational models of creativity. *Handbook of Creativity*: 351–373.

Carpuat, M., Ngai, G., Fung, P., and Church, K.W. (2002). Creating a Bilingual Ontology: A Corpus-Based Approach for Aligning WordNet and HowNet. *The Proceedings of GWC 2002, the 1st Global WordNet conference*, Mysore, India.

Dong, Z. (1988). Knowledge Description: What, How and Who? *The Proceedings of* *the International Symposium on Electronic Dictionaries*, Tokyo, Japan.

Hutton, J. (1982). *Aristotle's Poetics*. Norton.

Falkenhainer, B., Forbus, K. and Gentner, D**.** (1989). Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence*, 41, pages 1-63.

Fass, D. (1988). An Account of Coherence, Semantic Relations, Metonymy, and Lexical Ambiguity Resolution. *In: S. Small, G. W. Cottrell and M. K. Tanenhaus (eds.): Lexical Ambiguity Resolution: Perspectives from Psycholinguistics, Neuropsychology and A.I.* Morgan Kauffman: CA.

Kuhn, T. (1962). The *Structure of Scientific Revolutions*. Chicago: University of Chicago Press. J.

McCarthy, J. (1999). Creative Solutions to Problems. In the proceedings of the AISB'99 symposium on AI and Scientific Creativity (Edinburgh Scotland, 1999).

Miller, G. (1995). WordNet: A Lexical Database for English. *Comms. of the ACM*, 38(11).

Ritchie, G. (1999). Developing the Incongruity-Resolution Theory. *In the proceedings of the 1999 AISB Symposium on Creative Language: Stories and Humour*, Edinburgh, Scotland.

Veale, T. and Keane, M.T. (1997). The Competence of Sub-Optimal Structure Mapping on Hard Analogies. *The proceedings of IJCAI'97, the International Joint Conference on Artificial Intelligence,* Nagoya, Japan.

Veale. T. (2003). Dynamic Type Creation in Metaphor Interpretation and Analogical Reasoning: A Case-Study with WordNet. *The proceedings of ICCS2003, the 2003*

International Conference on Conceptual Structures, Dresden, Germany

Veale, T. (2004a). Incongruity in Humor: Root-Cause or Epiphenomenon? *The International Journal of Humor* 17/4, a Festschrift for Victor Raskin

Veale, T. (2004b). WordNet sits the S.A.T.: A Knowledge-Based Approach to Lexical Analogy. *The proceedings of ECAI'2004, the 16th European Conf. on Artificial Intelligence*. John Wiley: London.

Way, E. C. (1991). Knowledge Representation and Metaphor. *Studies in Cognitive systems*. Kluwer Academic Publishers: Holland.

Wong, S.H.S. (2004). Fighting Arbitrariness in WordNet-like Lexical Databases – A Natural Language Motivated Remedy. *The proceedings of GWC 2004, the 2nd Global WordNet conference*, Sojka, Pala, Smrz, Fellbaum, Vossen (eds.).

Context-Sensitive Category Structure in Lexical Ontologies

Tony Veale, Yanfen Hao,

School of Computer Science and Informatics, University College Dublin.

1. Introduction

Different contexts encourage different ways of speaking. This variation comprises more than differences in terminology and vocabulary — perhaps the most obvious reason for designing context-sensitive ontologies — but more insidiously comprises subtle differences in how common terms and categories are employed. Indeed, this variation in how ontological categories are expressed linguistically in different contexts often gives rise to domain-defining shibboleths; the plural term "ontologies", for instance, is more likely to identify its user as a computer scientist than as a philosopher (Guarino, 1998).

An ontology is a formalized and highly structured system of categories in which the meanings of semantic structures can be grounded. Guarino (1995) notes that such an “an engineering artifact, constituted by a specific vocabulary” is used to describe “a certain reality”, and so one expects this system to be fairly stable if it is to serve as a reliable bedrock of meaning. However, categories are no more than perspectives on the world, and these perspectives can change from context to context. For instance, when speaking of man-made objects, one can distinguish between the perspectives of designed functionality and ad-hoc functionality (see Barsalou, 1983). Banks do not design credit-cards so that they may be surreptitiously used to open locks, but in the context of certain movies and genres of fiction, this is an apparently frequent usage. Likewise, lamp-stands are not designed to be used as blunt instruments, or dinner plates as projectiles, yet these can be contextually appropriate functions for such objects. Clearly then, the categorization of an object depends not on the intrinsic type of the object as specified a priori in an ontology (though this can obviously be constraining), but on how the object is perceived in a particular context relative to a particular goal.

Once we accept that categorization is sensitive to context, all cognitive decisions that

follow from categorization — such as the perception of similarity between entities — become context-sensitive also. We see this effect in experiments performed by Morris (2006), which reveal that, in the context of an article about the effects of movies on suggestible teenagers, subjects reported a stronger semantic relationship between the terms "sex", "drinking" and "drag-racing". The context in question served to highlight the danger inherent in each of these activities, prompting the subjects to lump these ideas together under an ad-hoc and highly context-sensitive category of "dangerous behaviors" (see also Barsalou, 1983; Lakoff, 1987; Budanitsky and Hirst, 2006). This category is context-sensitive insofar as the specific behaviors that populate it are contextually determined. Smoking, for instance, is often considered a dangerous behavior in a medical context, but hardly seems to meet the diagnostic criteria for this category when viewed from the contexts of bomb-disposal, undercover police-work or high-wire acrobatics. So while we can expect the category Dangerous Behavior to be meaningful in any context of human behavior, the extension of the category may vary drastically from one context to another. A high-level division of labor between ontologies and contexts thus suggests itself: ontologies provide intensional definitions for categories that are meaningful in many contexts, while contexts provide extensional support for these categories in specific frames of reference or domains of application. As such, we see contexts and ontologies as comprising two complementary pieces of the larger knowledge-representation puzzle, a view consistent with that of Giunchiglia (1993) and Ushold (2000).

This work is thus a computational exploration of the common intuition that language use reflects category structure. As noted by De Leenheer and de Moor (2005), ontologies are, in the end, lexical representations of concepts, so we should expect that the effects of context on language use will 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 ontologies that naturally adapt to their contexts of use. Given this linguistic bias, we focus our attention in this paper to the class of ontology known as *lexical ontologies*. These are ontologies like WordNet (Fellbaum, 1998), HowNet (Dong and Dong, 2006) and the Generative Lexicon (Pustejovsky, 1995) that aim to serve as a formal ontological basis for a lexical semantics by combining knowledge of words with knowledge of the world.

Since many words and word-senses are inherently suited to some contexts of use more than others, the problem of context is one of particular importance to the proper working of such ontologies. Our focus on WordNet-like ontologies, lightweight as they are, is largely motivated by the fact that these ontologies have hitherto ignored the role of context in their design.

1.1. Structure of this Paper

We begin in section 2 by considering the interlocking roles of contexts and ontologies. We view each as a complementary kind of knowledge-representation, the primary distinction being one of stability: an ontology is a formal representation of concepts and their inter-relationships that is stable across different frames of reference, while a context is a changeable set of category-membership mappings from an ontology to the particulars of a given reference frame. The problem of "contextualizing an ontology" (Bouquet *et al.,* 2003; Obrst and Nichols, 2005) is thus seen as one of local categorization, in which categories with context-independent definitions are populated with context-specific members and membership-scores. In section 3 we describe how this contextualization can be computationally realized for a lexical ontology, not by modeling contexts directly and explicitly, but by using representative text corpora as sources of indicative linguistic behavior. These corpora yield local knowledge in the form of syntagmatic patterns, whereby, for instance, the patterns "X-addicted", "X-addled" and "X-crazed" suggest that the entity X is a kind of drug, the pattern "X-wielding" suggests that X is a kind of weapon, and "barrage of X" suggests that X is a kind of projectile. In section 4 we describe how stable ontological definitions can be automatically constructed from syntagmatic associations that are distilled from patterns of textual data on the web, in an approach that extends that of Almuhareb and Poesio (2005). We evaluate the reliability of these efforts in section 5, before concluding with some final remarks in section 6.

2. Context and Ontologies

As with any modeling task, ontological description is as much a matter of representational choice as it is one of representational verisimilitude. An ontology (qua

engineering artifact) does not capture "objective reality", or even a small portion thereof, but merely, as Guarino (1995) is careful to point out, "a certain reality". While a plurality of "realities" may be as confounding as a plurality of "ontologies", the term "reality" is nonetheless appropriate insofar as an ontology is designed to encode a common world-view that is shared by multiple (if not all) parties (see Guarino, 1998; Patel-Schneider *et al.,* 2003). The representational choice inherent in ontological design reflects the wide range of perspectives, biases, levels of detail and subject-oriented divisions that are available (consciously or otherwise) to the knowledge engineer. Regardless of the particular label one uses to motivate these choices, the notion of "context" seems to play the key role in defining the particular realities of different ontologies (e.g., see Bouquet *et al.,* 2003).

In distinguishing between ontologies and contexts, the former is often conceived as an inherently stable world-view, while the latter is conceived as an altogether more fluid and changeable frame of reference in which the former is applied. For instance, Obrst and Nichols (2005) conceive of contexts as user-dependent and task-dependent views on an underlying ontology, while Bouquet *et al.,* (2003) similarly conceive of contexts as local and private (i.e., unshared) perspectives onto a shared encoding of a domain. One role of context is to provide an additional layer of knowledge that better informs how an ontology can be used in a given set of circumstances. In particular, Obrst and Nichols suggest that context can serve to annotate or label the shared concepts and relations of an underlying ontology, to e.g., express the security-level and provenance of those elements. Imagine a medical ontology of diseases and their symptoms that can be applied in different socio-economic and geographic contexts. One would not expect the conceptual structure of the ontology to change from one context to the next, but one *would* expect the relative likelihood of different diseases, and the diagnostic value of different symptoms and tests, to differ significantly from, say, a highly developed country in the first world and a poorly developed country in the third world. Different contextualizing annotations of the same medical concepts could achieve precisely this multi-faceted effect.

Another role of context is to separate those parts of a knowledge-representation that are mutually inconsistent into different, but complementary, perspectives, each perhaps owned by a different agent. As used in the Cyc ontology (Lenat and Guha, 1990), these perspectives are called microtheories; their propositional content remains local and private unless explicitly inherited by other microtheories or made visible through a process of lifting (Guha, 1991). For instance, the concept Sherlock-Holmes can be ontologized as a kind of fictional character, and thus, a kind of mental product, or it might be ontologized as a kind of detective. WordNet opts for the former course rather than the latter, thus sacrificing the ability to reason about Holmes as if he were an actual detective, or even an actual person. In an ideal ontology, both ontological perspectives would be made available for reasoning purposes, perhaps by representing each in a separate microtheory, or by representing each in different ontologies and providing a detailed system of mappings between each (e.g., as in Bouquet *et al.,* 2003).

Each of these apparent roles sees context as a means of partitioning ontological content into alternate views of reality. Indeed, the microtheory labels used by Cyc, such as HealthMt, HistoryMt and so on, can be seen as annotations on propositions that allow Cyc's inference processes to selectively include or exclude large swathes of the ontology's content in a given reasoning process. In this vein, another related role of context is to provide a bridge between the stable definitions of an ontology and the contingent facts of a particular world-view. For instance, an ontology of chemical substances may be agnostic with respect to how those substances are used, so that the same substance might be categorized as a medicine in one context, an illegal drug in another, and a poison in yet another. It is this division of labor between ontologies and contexts that interests us most in this current work: how can we create ontologies as collections of stable category definitions that apply in all contexts, yet which are instantiated differently, by different entities and to different degrees, in specific contexts? Given the significant design and engineering efforts that are employed in the construction of well-formed ontologies (e.g., see Gangemi *et al.,* 2001), this division of labor should be a clean one, so that the base ontology only posits relationships that are safe in all contexts, and each context only posits relationships that complement, rather than contradict, those of the base ontology.

This division of labor requires a solution to two related computational problems: how do we acquire and represent the stable category definitions that comprise the ontology;

and how do we acquire the local contextual distinctions that cause these definitions to be instantiated by different entities in different frames of reference? Almuhareb and Poesio (2005) describe a web-based approach to acquiring the attribute structure of concepts via text analysis of internet content, as indexed by a search engine like Google. Their approach indicates how both stable category structures and contingent instantiations of those structures can be inferred from simple processes of text analysis. Almuhareb and Poesio use highly diagnostic search queries such as "*the * of a|an|the C is|was*" to identify attributes of a given concept C in web texts. By acquiring attributes (such as the fact that beverages have an associated temperature and strength) as opposed to simple attribute values (such as "hot" and "cold" for coffee), these authors acquire a general frame structure for each concept that can be instantiated differently in different contexts.

We too employ a large-scale analysis of web-text to acquire stable category definitions that will transcend context boundaries. However, we do not currently focus on the acquisition of attribute structure, but on prototypical attribute values. While Almuhareb and Poesio (2005) demonstrate that generic attributes such as Temperature and Colour are more revealing about conceptual structure that specific values such as "hot" and "red" (since these values can change without affecting the nature of the category), we do not collect arbitrary contingent attributions (such as the fact that coffee *can* be cold) but highly diagnostic and category-defining attributions (e.g., that espresso *should* be strong, that surgeons *should* be delicate, that gurus *should* be wise, and so on). To identify which attribute values are truly central to the consensus definition of a category, we use the highly specific comparison frame "as * as a|an C" to collect similes involving a given category from the web. Once acquired and validated, we articulate the prototypical attributes for a given category as a set of logical constraints that serves as a functional definition for that category. The form of these functions is presented in section 3, while the web-based acquisition of each function's content is described in section 4.

2.1. Texts and Con-Texts

De Leenheer and de Moor (2005) see a context as a mapping from a set of lexical (and potentially ambiguous) labels to a set of language-neutral concept identifiers. In this view, the same words can denote members of different categories in different contexts.

For instance, “cocaine” can denote a kind of local anesthetic in a medical context and a kind of illegal drug in a law-enforcement context. This is more than a matter of lexical ambiguity; the same sense is intended in each context (i.e., the same substance) but a different ontological categorization is implied in each. Combining a mapping-theoretic view of context (e.g., Obrst and Nichols, 2005; Bouquet *et al.,* 2003), with the lexical emphasis offered by De Leenheer and de Moor, it is possible to obtain much of the reasoning benefits of a context without an explicit logical representation of context. Different contexts give rise to different ways of speaking, from slang to specialist terminology to habitual turns of phrase, so that the localizing effect of context can often be inferred from a linguistic analysis of a representative text corpus (see Hoey, 2000; Pustejovsky, Hanks and Rumshisky, 2004). The most revealing analyses will be syntagmatic in nature, looking beyond particular word choices to larger patterns of contiguous usage. For example, the similarity between chocolate and a narcotic like heroin will, in most contexts, simply reflect the ontological fact that both are kinds of substances; certainly, taxonomic measures of similarity as discussed in Budanitsky and Hirst (2006) would capture little more than this basic categorization. However, in a context in which the addictive properties of chocolate are highly salient (in an on-line dieting forum, for instance), chocolate is more likely to be categorized as a drug and thus by 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 (each of these uses is to be found in chocolate-related Wikipedia articles). In a context that gives rise to these expressions, it is unsurprising that chocolate should appear altogether more similar to a harmful narcotic.

A given corpus may employ syntagmatic patterns which reflect the fact that the corresponding context views chocolate as a kind of drug, or military robots as soldiers, or certain kinds of criminal as predators. By augmenting a base ontology with these categorizations, the ontology may become sufficiently contextualized to reason fluently in this context. The model of corpus-based ontology augmentation we describe in this paper is consistent with, and complementary to, the Theory of Norms and Exploitations (TNE) proposed by Hanks (2004), in which corpus analysis is used to identify both the

syntagmatic *norms* of word usage (i.e., highly conventional and normative uses) and meaning-coercing *exploitations* of these norms.

2.1. Finding Ontological Insights in Text

A syntagmatic approach to deriving ontological insights from text is hardly novel. Hearst (1992) describes a syntagmatic technique for identifying hyponymy relations in free text by using frequently occurring genre-crossing patterns like NP_0 *such as* $\{NP_1, NP_2, \ldots, NP_n\}$. Like the approach of Charniak and Berland (1999), Hearst’s patterns exploit explicit mentions of category structure, as in the phrase “drugs like Prozac, Zoloft and Paxil”. Such techniques are useful because contexts frequently introduce new terms that are locally meaningful. Nonetheless, such techniques do not reveal the subtle differences in category usage that underpin a particular context. These differences are implicit precisely because the existence of a context presupposes the existence of a shared body of knowledge and a common world-view. Context-specific corpora only reveal this shared knowledge indirectly, insofar as it is presupposed in the way that language is used.

Closer to the current approach is that of Cimiano, Hotho and Staab (2005), who do not look for unambiguous “silver bullet” patterns in a text, but who instead characterize a lexical term according to the syntagmatic patterns in which it participates. These patterns include the use of the term as the subject, object or prepositional complement of a verb. The key intuition, expressed also in Weeds and Weir (2005), is that terms with similar distribution patterns will denote ideas that are themselves similar. Cimiano *et al.* exploit the phrasal dependencies of a term as features of that term that can be used, through a process of conceptual clustering called Formal Concept Analysis, to determine subsumption relations between different terms. At no point are explicit expressions of these relations sought in a text. Rather, from a tabular mapping of terms to their syntagmatic attributes (called a Formal Context), FCA is used to infer this relation by determining which terms possess attribute descriptions that are a superset or subset of other descriptions. These attributive descriptions serve a dual purpose: they allow an extensional comparison of different concepts to determine which is more general and inclusive; but they also serve as an explicit intensional representation of the conceptual terms that are ontologized. For example, the term “bike” is recognized as *rideable*, *bookable* and *rentable* because of its use as an object with the verbs “ride”, “book” and “rent”, so the set *(rideable, bookable, rentable}* provides an intensional picture of Bike.

Our approach likewise employs syntagmatic properties as semantic features over which ontological categories can be defined. Our focus, however, is on providing a functional description framework for ontologists, one that allows the intensional definition of a category to be expressed in a context-independent manner, yet which can be applied to a corpus to discover contextually appropriate instances of that category. A description of this framework is offered in the next section, before, in section 4, we describe a semi-automated means for constructing functional definitions for the most commonly used lexical concepts in the WordNet ontology,

3. Category Norms and Contextual Exploitations

In this section we present a functional framework for defining ontological categories in terms of how those categories are expected to behave when expressed linguistically. These expectations are articulated as syntagmatic norms (Hanks, 2004) that capture e.g., the most diagnostic adjectival modifiers that contribute to a lexical description of the category, the kinds of verbs for which the category typically acts as an agent or a patient, the kind of group terms (like “army”, “herd”, “flock”, etc.) that are typically used to describe aggregations of the category, and so on. Each ontological category is assigned a different functional form that expresses the appropriate syntagmatic expectations. This functional form serves as a membership function for the category, returning a membership score in the range 0 … 1, where 0 denotes no membership, and 1 denotes the membership level of the most prototypical and exemplary members.

A continuous or fuzzy membership score allows radial category structures (as described in Lakoff, 1987) to be defined and populated for a given context/corpus. For example, to the extent than a collocation like “army of X” is found in a corpus, the associated context can be said to categorize X as a sub-type of Soldier. Likewise, to the extent that the syntagm “X-addicted” has currency in a corpus, X should be seen as a kind of Drug. Interestingly, some of the most stable and unambiguous syntagmatic patterns are

associated with metaphoric conceptualizations. Thus, the syntagmatic schema “barrage of X” identifies X as a projectile, whether X is an arrow, a pointed question or an angry email. The frequency of these patterns in a corpus yields a sliding scale of category membership in the associated context. Thus, an entity can be more representative of a particular category in one context (say, Chocolate as a Narcotic in a weight-loss context) than in another.

We begin by supposing a function *(attr* arg_0 arg_1*)* that returns a real number in the range [0..1] based on the frequency of arg_0 as an adjectival modifier for the noun arg_1 in a corpus. Suppose also a function *(%isa* arg_0 arg_1*)* that returns a number in [0..1] reflecting the proportion of senses of arg_0 that are descendants of arg_1 in a base-ontology like WordNet. We can now define the category Fundamentalist in a functional fashion:

(define Fundamentalist (arg_0*)*
    *(**   *(max (%isa* arg_0 *Person) (%isa* arg_0 *Group))*
        (min
            (max (attr political arg_0*) (attr religious* arg_0*))*
            (max (attr extreme arg_0*) (attr violent* arg_0*) (attr radical* arg_0*)))*
)
)

Figure 1: A functional description of the category Fundamentalist

That is, any extreme, violent or radical person or group that is either political or religious deserves to be categorized as a fundamentalist. The extent to which this person or group is a fundamentalist depends entirely on the contextual evidence for these criteria, as measured by the function *attr*. The precise workings of *attr* can be implemented in a number of ways, using any of a variety of corpus-based distributional similarity metrics, such as Dice’s coefficient or the Jaccard measure (see Lee, 1999; Weeds and Weir, 2005). Whatever measure is used, it must either return a value in the range [0..1] or be scaled to do so, so that each user-defined function like *Fundamentalist* will likewise return a value in the [0..1] range. The value returned by an intensional function thus corresponds to a membership score in the corresponding radial category (Lakoff, 1987). Thus, if *(Fundamentalist evangelical)* returns a score of 0.41 for a given corpus, this corresponds to a high membership in the category Fundamentalist. These user-defined membership functions also serve as explicit symbolic representations of intensional structure, inasmuch as they can be given a logical interpretation. In the example of Fundamentalist, note how the mathematical functions *min* and * (multiplication) are essentially used to encode a fuzzy-logic equivalent of the logical operator *and*, while the function *max* encodes a fuzzy-logic equivalent of the logical operator *or*.

Table 1: *Basic Category-defining functions and their syntagmatic correspondences*

Function	Example	Syntagmatic Pattern(s)	Range
(agent $verb_0$ $noun_0$*)*	(agent kill robot)	"$noun_0$ $verb_0$" "… $verb_0$+past by $noun_0$"	[0 … 1]
(patient $verb_0$ $noun_0$*)*	(patient eat prey)	"$noun_0$ verb+past by …" "… $verb_0$ $noun_0$"	[0 … 1]
(attr adj_0 $noun_0$*)*	(attr knight brave)	"adj_0 $noun_0$" "as adj_0 as a | an $noun_0$"	[0 … 1]
(group $noun_0$ $noun_1$*)*	(group army grunt)	"$noun_0$ of $noun_1$+plural"	[0 … 1]
(of $noun_0$ $noun_1$*)*	(of owner pet)	"$noun_0$ of $noun_1$"	[0 … 1]
(hyphen $verb_0$ $noun_0$*)*	(hyphen shape egg)	"$noun_0$- $verb_0$+past" (e.g., egg-shaped, bite-sized)	[0 … 1]

Category-membership functions can exploit most of the syntagmatic patterns employed in Cimiano *et al*. (2005), with some additions (see Table 1). For instance, the "GROUP of NOUN+plural" pattern employs WordNet to identify group membership descriptions in a corpus, where GROUP is any group-denoting WordNet term (e.g., swarm, army) or group activity (e.g., barrage, invasion, influx). This syntagm is exploited in intensional descriptions via the function *(group* arg_0 arg_1*)*, which returns the extent (in the range [0..1]) to which arg_1 is described as a member of the group arg_0 in a given corpus. For instance, using the text of the encyclopaedia Wikipedia as a corpus, we find *(group influx immigrant) = 0.38* and *(group army mercenary) = 0.31* (using Dice's coefficient as a measure of salience). The base functions of Table 1 thus serve as the interface between a

context-independent intensional description, like that of Fundamentalist in Figure 1, and a specific context-defining corpus. Our intension for Fundamentalist is thus context-neutral, but can be applied to any corpus to yield a context-specific extension.

Some syntagmatic patterns are more obvious about category membership than others. For instance, a word form like "mint-flavored" indicates that Mint is a flavor. Hyphenated forms can also be used to indicate figurative membership in categories like:

(define Drug(arg_0*) (hyphen addict* arg_0*))* *% e.g., risk-addicted*

(define Causal-Agent(arg_0*) (hyphen induce* arg_0*))* *% e.g., drug-induced*

It is also useful to view the intensional description in Figure 1 as a structured-query in an information retrieval system. In this case, the query is designed to retrieve not documents, but category members, while the corpus is indexed not on keywords, but on meaning-rich syntagmatic patterns. After all, IR queries essentially define ad-hoc categories (Barsalou, 1983) whose members are documents and whose degrees of membership are given as relevance scores. Just as a retrieval query marshals evidence for a given document using weighted query operators, each category membership function marshals syntagmatic evidence that will yield a final membership score. To this end, we introduce another basic function, *combine*, which will allow us to combine multiple pieces of evidence into a score in the [0..1] range. If e_0, e_1, etc. are the scores associated with various pieces of evidence (as returned by the functions of table 1, say), then *combine* adds these scores to yield another in the [0..1] range thus:

$$(\text{combine } e_0\ e_1) = e_0 + e_1(1 - e_0) = e_0 + e_1 - e_0 e_1$$

$$(\text{combine } e_0\ e_1 \ldots e_n) = (\text{combine } e_0\ (\text{combine } e_1 \ldots e_n))$$

This *combine* function is thus a naïve probabilistic *or* function, one that naively assumes independence among the evidence it combines to generate scores that asymptotically approach 1.0. If a piece of evidence is included multiple times (to increase its relative contribution), it is counted multiple times, but with a diminishing effect.

Consider the use of *combine* in a category definition for Invader in Figure 2. Note how four types of information are synthesized in this definition: general taxonomic knowledge (via the *%isa* function); adjectival modification (via the *attr* function); subject-verb knowledge (via *agent*); and group membership knowledge (via *group*).

```
(define Invader (arg0)
      (combine    (* 0.3 (max (%isa arg0 Person)  (%isa arg0 Group)))
                  (agent invade arg0)
                  (attr invasive arg0)
                  (group invasion arg0)
                  (group influx arg0)
                  >=2
      )
)
```

Figure 2: A functional description of the category Invader

The final clause *>=2* in Figure 2 is called a "quantitative cut": it specifies the number of non-zero arguments that *combine* must have processed prior to this cut if it is to perform its normal function; if this threshold is not met, then *combine* aborts (i.e., cuts) early and simply returns a 0. Therefore, any term in a given context that meets two or more of these intensional criteria (e.g., people or groups that invade, non-human invasive organisms that form an influx, etc.) is categorized as an Invader with a degree of membership that is a function of the evidence available in the corpus. Note how the contribution of WordNet (or whatever ontology underpins the *%isa* function) is scaled by a small multiplier of 0.3. This prevents the *%isa* clause – which merely serves as a kind of soft-preference or constraint here – from making an undue contribution to the overall membership score.

3.1. Introspection about Category Membership

Consider a membership function for the category Pet which, as formulated in Figure 3, combines several different types of evidence to diagnose "pet-hood". The definition asks the following questions of is each potential member: is it a kind of animal? Is it docile or domesticated? Is it cute? Is it something that one can own and care for? For those terms that contextually meet two or more of these demands in a context/corpus, this definition can be used to introspectively explain why.


```
(define Pet (arg0)
      (combine    (* 0.3  (%isa arg0 Animal))
                  (max (of owner arg0)    (of care arg0))
                  (max (attr docile arg0) (attr domesticated arg0))
                  (max (attr cute arg0)    (attr cuddly arg0))
                  >=2
      )
)
```

Figure 3: A functional description of the highly context-sensitive category Pet

The definition of Figure 3 is also constructed to make membership in a normative class (here Animal) a soft-preference rather than a hard constraint, since one can conceive of human pets (favored children, slaves) and artificial pets (toys, robots, etc.) and so on. Suppose, in a given context, the above function assigns a membership of 0.12 to the term Iguana. Introspecting over the symbolic structure of the definition, the system can explain why this assignment was made, by pointing out that the associated corpus speaks of iguanas as cuddly or cute with this much frequency, and as docile with that much frequency, and so on. Now suppose a zero membership is given to Piranha. The system can use a similar process to perform a what-if analysis, much like a spreadsheet. Looking at the ancestry of Piranha in a base ontology like WordNet, the system can determine which of the elements in the definition are applicable to Piranha. Noting from the corpus that cuddliness, cuteness and docility are collocates of "animal", it can then explain that Piranha is not a Pet because it is seen as neither cute, cuddly or docile in this context.

4. Web-based Acquisition of Category-Membership Functions

The functions of Figures 1,2 and 3 make no reference to any kind of context. Rather, they encode diagnostic knowledge of a general character about individual categories – what one might describe as the conventional wisdom about these categories. Our definitions of Pet, Invader, Fundamentalist, Drug, etc. are intended to represent quantitative membership functions for the correspondingly-named categories in a broad-scope lexical ontology like WordNet. They should thus be seen not as comprising a local ontology of their own, but as additions to a base ontology (see Giunchiglia, 1993; Ushold, 2000).

Nonetheless, these functions are inherently context-sensitive, in two crucial respects. Firstly, they encode conventional wisdom about category structure in a flexible manner, not as hard constraints but as soft preferences. In this way, they anticipate that certain contexts may observe certain diagnostic requirements and not others, e.g., that invaders are not always human, or that pets may not always be animals.

Conventional wisdom has its own syntagmatic norms of expression. For instance, when one wishes to highlight a specific property in a given entity, it is commonplace to compare that entity to one for which that property is widely agreed to be diagnostic. Comparisons of the form "as ADJ as a/an NOUN" work best when the exemplar that is used (e.g., "dry as *sand*", "hot as the *sun*") is familiar to the target audience and is truly exemplary of the given property in a context-independent manner. That is, such simile-based comparisons work best when they are generally self-evident and not dependent on a private context to give them meaning. By searching the web for comparisons of this form, we achieve two important results: we identify the exemplar concepts that are most frequently used as a basis of comparison, and which are thus most deserving of ontological representation; and, we identify the most salient properties of those exemplars categories, allowing us to compose category membership functions for them.

As in Almuhareb and Poesio (2005), we use the Google API to find instances of our search patterns on the web. We use two simile patterns, one in which the wildcard operator * substitutes for the adjectival property (where an exemplar noun is explicitly given), and one in which the wildcard operator substitutes for the noun (while the adjective is given). The first pattern collects salient adjectival properties for a given noun, while the second collects the most common noun concepts that exemplify a given adjectival property. For purposes of radial category construction, we expect that adjectives which denote a point on a scale, such as "brave" (versus "cowardly"), "hot" (versus "cold") and "rich" (versus "poor") will be the most commonly used adjectives in comparative phrases, and will yield the most diagnostic features for categorization. We initially limit our attention then to WordNet adjectives that are defined relative to an antonymous term. For every adjective ADJ on this list, the query "*as ADJ as **" is sent to Google and the first 200 snippets returned are scanned to extract different noun bindings

(and their relative frequencies) for the wildcard *. 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 hinge on non-antonymous adjectives that were 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.

4.1. Feature Filtering

The simile frame "as ADJ as a NOUN" is relatively unambiguous as such patterns go, but a non-trivial number of unwanted or noisy data is nonetheless retrieved. In some cases, the NOUN value forms part of a larger noun phrase that is not lexicalized in WordNet: it may be the modifier of a compound noun (e.g., "bread lover"), or the head of complex noun phrase (such as "gang of thieves" or "wound that refuses to heal"). In other cases, the association between ADJ and NOUN is simply too ephemeral or under-specified to function well in the null context of a base ontology. As a general rule, if one must read the original document to make sense of the association, 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 of the ironies we found exploit contingent 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 human judges, 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 judges also annotated 4685 simile instances (of 2798 types) as ironic; these similes relate a set of 936 adjectives to a set of 1417 nouns. Surprisingly, ironic pairings account for over 13% of all annotated simile instances and over 20% of all annotated simile types.

4.2. Linking to WordNet

WordNet is used as a source for the adjectives that drive the simile retrieval process; it is also used to validate the nouns (unitary or multi-word) that are described by these similes. By sense-disambiguating these nouns relative to the noun-senses found in WordNet, we can use their associated adjectival features to construct functional category definitions (as described in section 3, and illustrated in Figures 1,2 and 3) for each of these WordNet senses. That is, we can automatically construct a set of context-sensitive membership functions for the most commonly used terms in the WordNet noun ontology.

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. Two similes “as A_0 as N_1” and “as A_0 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 noun is likely to be

sufficiently exemplary of some property to be useful in a simile.

4.3. From Similes to Membership Functions

The above filtering and WSD processes associate the features *stealthy*, *silent* and *agile* with the person sense of “ninja” (denoted ninja.0), leading to the following function:

```
(define Ninja.0 (arg0)
      (*    (%isa arg0 Person.0)
            (combine      (attr  stealthy arg0)
                          (attr  silent arg0)
                          (attr  agile arg0)
                          >=2)
      )
)
```

Figure 4: A web-derived description of the noun “ninja”

As we cannot know which subset of these features is sufficient for categorization, we use the quantitative cut *>=2* to ensure that more than feature is contextually present to support a categorization as a ninja. The more features that are present, the higher the resulting membership score (aggregated via the *combine* operator) will be. The factor *(%isa* arg_0 *Person.0)* is chosen as the taxonomic constraint for all category functions that represent a specialized kind of person in WordNet.

The most commonly used bases of comparison will provide more features to choose from, and will thus yield a more finely discriminated range of members across different contexts. Consider the membership function for Snake as illustrated in Figure 5.

```
(define Snake.0 (arg0)
   (*    (%isa  arg0  Animal.0)
         (combine  (attr cunning arg0)      (attr slippery arg0)
                   (attr slim  arg0)        (attr flexible arg0)
                   (attr sinuous  arg0)     (attr crooked arg0)
                   (attr deadly  arg0)      (attr poised arg0)
                   >= 2)))
```

Figure 5: A web-derived description of the animal sense of “snake” (snake.0)

Note that the taxonomic constraints *(%isa* $\arg_0$ *Person.0)* and *(%isa* $\arg_0$ *Animal.0)* serve to ensure that the resulting in-context categorizations are broadly literal w.r.t. WordNet. By weakening (i.e., generalizing) or removing these constraints, one could allow for contextually-appropriate metaphoric categorizations to made, e.g., that agile animals or stealthy and silent organizations might be seen as ninjas, or that cunning and slippery people might be seen as snakes. The distinction between literal and metaphoric categorization in a given context is often blurred, and may, in principle, be impossible to delineate. Is chocolate really an addictive drug in some dieting contexts, or is such a categorization simply a figurative stretch? While this rather vexing question falls outside the scope of the current paper, we note that the framework of category membership functions described here provides an ideal mechanism for exploring the contextual boundaries of literal and metaphoric categorization in future research.

5. Empirical Evaluation

In this section we provide empirical support for the two main claims of this paper. The first is the relatively uncontroversial claim that syntagmatic patterns of usage at the word-level reflect distinctions in category usage at the ontological level (e.g., Cimiano *et al.*, 2005; Hanks, 2006), so that the syntagmatic patterns of a given corpus can be taken to be indicative of category membership in the corresponding context. The second is the more novel claim that similes are sufficiently revealing about the diagnostic properties of categories to allow accurate category membership functions to be constructed.

We test the first claim using the HowNet ontology of Dong and Dong (2006), HowNet differs from WordNet in many respects (e.g., the former is bilingual, linking the same definitions to both English and Chinese labels) but the key difference is that HowNet defines the meaning of each word sense via a simple conceptual graph. For instance, HowNet specifies that a Knight is the agent of the activity Fight, while Assassin is the agent of the activity Kill. Additionally, it states (in explicit logical terms) that the killing performed by an Assassin has the attribute *means=unlawful*. Each of these logical definitions is hand-crafted, allowing us to test whether the syntagmatic patterns in a corpus would suggest the same semantic structures as a trained knowledge-engineer. For

simplicity, we focus on those categories that are defined as an agent of a particular activity, like Knight and Assassin. Using the complete text of Wikipedia as our corpus (2 gigabytes, from a June 2005 download), this corpus contains 1626 different nouns that have at least one sense that HowNet defines as filling the agent case of a specific activity. In all, HowNet uses 262 unique verbs, like *kill*, *buy* and *repair* to describe these activities. Using Dice's coefficient (Lee, 1999) to measure the association between each noun and each verb for which the noun is used as an active subject, we find that in 69% of cases, the highest rating is given to the verb that HowNet itself uses to encode the noun's sense.

Our second claim concerns the simile-gathering process of the last section, which, abetted by Google's practice of ranking pages according to popularity, should reveal the most frequently-used nouns in comparisons on the web, and thus, the most useful categories to annotate in a lexical 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 used 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. One measurable clue as to how a category is perceived is its affective rating.

For instance, humans – unlike computers – tend to associate certain positive or negative feelings, or affective values, with particular categories. 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 web-derived simile data). More specifically, we should expect that ratings estimated via these simile-derived properties to exhibit a higher correlation with the independent ratings of Whissell's dictionary than properties derived from other sources (such as WordNet itself) or from other syntagmatic frames.

To determine if this is indeed the case, we calculate and compare this correlation between predicted and reported affect-ratings using the following data sources:

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 the shallow-parsing of WordNet glosses; e.g., *strong* and *black* are extracted from the gloss for Espresso ("strong black coffee brewed by forcing steam under pressure …").

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 ($p < 0.05$). As expected, attribute values derived from bona-fide similes only (A) yielded the best correlation (+0.5) while attribute values derived from ironic similes only (C) yielded the worst (-0.2); a middling correlation coefficient of 0.32 was found for all similes together (B), reflecting the fact that bona-fide similes outnumber ironic similes by a ratio of 4 to 1. A weaker correlation of 0.25 was found using the corpus-derived adjectival modifiers for each noun (D); while this data provides quite large value sets for each noun, these attribute values merely reflect the potential rather than intrinsic properties of each category and so do not reveal what is most diagnostic about the category. As also noted by Almuhareb and Poesio (2005), such values reveal very little about the conceptual structure of a category. Those authors address this problem by instead seeking to mine the attributes (i.e., frame-slots) rather than their potential values, while we address the problem by only mining the most diagnostic attribute values.

More surprisingly, perhaps, property sets derived from WordNet glosses (E) are also poorly predictive, yielding a correlation with Whissell's affect ratings of just 0.278. Our

goal in this paper has been to describe a framework for augmenting WordNet's categories with membership functions that both reflect the diagnostic properties of these categories and that allow them to apply to different entities in different contexts. These results suggest that the properties needed to construct these membership functions are not to be found within WordNet itself, but must be acquired by observing how people actually use categories to construct and convey meanings.

6. Concluding Remarks

In this paper we have described a category representation that serves as a flexible interface between, on one hand, the need for ontological clarity and commitment to explicit category definitions, and on the other, the context-sensitive utilization of these definitions. Those semantic properties that are diagnostic of category-membership are given explicit intensional form in category definitions that can also be used as structured queries that construct their own context-sensitive extensions from a representative text corpus. These category definitions, essentially category-membership functions, establish their own boundaries based on the context, and – under the ontologist's control – can blur the traditional line between literal and metaphoric category usage when it is ontologically useful to do so (e.g., see Hanks, 2006).

This programmatic approach to category definition complements the syntagmatic approach to ontology construction outlined in Cimiano *et al.* (2005), whereby the ontologist is given access to syntagmatic features of a context via a flexible but powerful representation language. Nonetheless, we have also described how category definitions can, like those of Cimiano *et al*., be created automatically. The key to this automation is the identification of the most salient and diagnostic attribute values for the members of a category by analysis of the most frequently used comparative phrases on the web.

By constructing functional category definitions for the most commonly used WordNet senses, we achieve a pair of related goals: WordNet is augmented with a robust, non-classical view of category structure; and, perhaps more importantly (in the context of this special issue, at least), WordNet is remade in a context-sensitive form. Ultimately, these two goals are flip-sides of the same coin, for insofar as context alters the perceived boundaries of familiar categories, classically-structured ontologies cannot be made context-sensitive without first being augmented with a flexible sense of category membership.

Much work remains to be done on the current framework, not least on the tantalizing issue of where literal categorization ends and metaphoric categorization begins, and the role of context in blurring this boundary. While our primary focus in this paper has been on the definition of categories in terms of adjectival modification and the *attr* syntagmatic pattern, other productive patterns abound, as described in section 3. In this vein, we are currently employing the simile-gathering approach to acquire the most salient behaviors of categories, by using the comparison frame “to VERB like a/an NOUN”. By marshalling as wide a range of syntagmatic insights as possible in the construction of category membership functions, our hope is that lexical ontologies like WordNet can more robustly meet the challenges of a changing context while remaining – in the sense of Gangemi *et al.* (2001) – as ontologically “clean” as possible,

References

Almuhareb, Abdulrahman and Poesio, Massimo. 2005 Concept Learning and Categorization from the Web. In *Proceedings of CogSci 2005, the 27th Annual Conference of the Cognitive Science Society*. New Jersey: Lawrence Erlbaum.

Bouquet, Paolo, Giunchiglia, Fausto, van Harmelen, Frank, Serafini, Luciano, and Stuckenschmidt, Heiner. 2003. C-OWL: Contextualizing Ontologies. In Proc. 2nd International Semantic Web Conference, LNCS vol. 2870:164-179. Springer Verlag.

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

Barsalou, Lawrence W. 1983. Ad hoc categories. Memory and Cognition, 11:211–227.

Charniak, Eugene and Berland, Mathew. 1999. Finding parts in very large corpora. In *Proc. of the 37th Annual Meeting of the Assoc. for Computational Linguistics,* pp. 57–64.

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

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

Dong, Zhendong and Dong, Qiang. 2006. *HowNet and the Computation of Meaning*. World Scientific: Singapore.

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

Gangemi, A., Guarino, N., and Oltramari, A. 2001. Conceptual Analysis of Lexical Taxonomies: The Case of WordNet’s Top-Level. In C. Welty and S. Barry (eds.), *Formal Ontology in Information Systems*. *Proceedings of FOIS2001*. ACM Press: 285-296.

Guarino, Nicola and Giaretta, Pierdaniele. 1995. Ontologies and Knowledge Bases. Towards a terminological clarification. In N. J. I. Mars (ed.), Towards Very Large Knowledge Bases. Amsterdam: IOS Press.

Guarino, Nicola. (ed.) 1998. Formal Ontology and Information Systems. Amsterdam: IOS Press. . *Proceedings of FOIS1998*, June 6-8, Trento, Italy.

Giunchiglia, Fausto. 1993. Contextual reasoning. *Special issue on I Linguaggi e le Macchine* XVI:345–364.

Gruber, Thomas. 1993. A translation approach to portable ontologies. *Knowledge Acquisition*, 5(2):199-220.

Guha, Ramanathan, V. 1991. Contexts: a formalization and some applications. *Technical Report STAN-CS-91-1399*, Stanford Computer Science Dept., Stanford, California.

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

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

Hearst, Marti. 1992. Automatic acquisition of hyponyms from large text corpora. In *Proc. of the 14th International Conference on Computational Linguistics*, pp 539–545.

Hoey, Michael. 2000. *Patterns of Lexis in Text*. Oxford University Press.

Lee, Lillian. 1999. Measures of distributional similarity. In *Proc. of the 37th Annual Meeting of the Association for Computational Linguistics,* pp. 25–32.

Lenat, Douglas and Guha, Ramanathan, V. 1990. *Building Large Knowledge-based Systems: Representation and Inference in the CYC project*. Addison Wesley.

Obrst, Leo and Nichols, Deborah. 2005. Context and Ontologies: Contextual Indexing of Ontological Expressions. *In Proc. of the AAAI 2005 Workshop on Context and Ontologies, Pittsburgh, Pennsylvania*.

Pustejovsky, James. 1995. Generative Lexicon. The MIT Press, Cambridge, MA.

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

Lakoff, George. 1987. *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*. University of Chicago Press, Chicago.

Morris, Jane. 2006. Readers' Perceptions of Lexical Cohesion and Lexical Semantic Relations in Text. *Ph.D. thesis*, University of Toronto.

Patel-Schneider, Peter F., Hayes, Patrick and Horrocks, Ian. 2003. Web Ontology Language (OWL) Abstract Syntax and Semantics. Technical report, W3C, www.w3.org/TR/owl-semantics.

Ushold, Michael. 2000. Creating, integrating and maintaining local and global ontologies, In *Proc. of the 1st Workshop on Ontology Learning (OL-2000), part of* ECAI 2000.

Ushold, Michael and Gruninger, Michael. 1996. Ontologies: Principles, methods and applications. *The Knowledge Engineering Review* 11(2):93–136.

Weeds, Julie and Weir, David. 2005. Co-occurrence retrieval: A flexible framework for lexical distributional similarity. *Computational Linguistics*, 31(4):433–475.

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

Tracking the Lexical *Zeitgeist* with WordNet and Wikipedia

Tony Veale[1]

Abstract. Most new words, or neologisms, bubble beneath the surface of widespread usage for some time, perhaps even years, before gaining acceptance in conventional print dictionaries [1]. A shorter, yet still significant, delay is also evident in the life-cycle of NLP-oriented lexical resources like WordNet [2]. A more topical lexical resource is Wikipedia [3], an open-source community-maintained encyclopedia whose headwords reflect the many new words that gain recognition in a particular linguistic sub-culture. In this paper we describe the principles behind *Zeitgeist*, a system for dynamic lexicon growth that harvests and semantically analyses new lexical forms from Wikipedia, to automatically enrich WordNet as these new word forms are minted. *Zeitgeist* demonstrates good results for composite words that exhibit a complex morphemic structure, such as portmanteau words and formal blends [4, 5].

1 INTRODUCTION

Language is a dynamic landscape in which words are not fixed landmarks, but unstable signposts that switch directions as archaic senses are lost and new, more topical senses, are gained. Frequently, entirely new lexical signposts are added as newly minted word-forms enter the language. Some of these new forms are cut from whole cloth and have their origins in creative writing, movies or games. But many are patchwork creations whose origins can be traced to a blend of existing word forms [1]. This latter form of neologism is of particular interest to the computational lexicographer, since such words possess an obviously compositional structure from which one can begin to infer meaning. In this paper, we demonstrate that, if given enough semantic context, an automated system can assign a sufficiently rich semantic structure to these words to allow them to be automatically added to an electronic database like WordNet [2]. When tied to a system for harvesting new word forms from the internet, this capability allows for a dynamic WordNet that grows itself in response to a changing language and cultural context.

Most neologisms bubble beneath the surface of widespread usage before they gain entry to a conventional dictionary. This is to be expected, since the internet is awash with idiosyncratic neologisms that lack both charm and staying power. Nonetheless, to experience the variety and inventiveness of the most creative new words in English, one need look no further than Wikipedia [3], an open-source electronic encyclopedia that is continuously updated by a on-line community of volunteers. If such words are likely to be encountered in any text to which NLP technologies are applied, from deep text understanding to shallow spell-checking, we should expect our lexical databases to possess a basic interpretation capability.

In this paper, we describe an automated system, called *Zeitgeist*, that harvests neologisms from Wikipedia and uses the semantic context provided by Wikipedia's topology of cross-references to add corresponding semantic entries to WordNet. In section two we briefly describe WordNet and Wikipedia, and outline the properties of each that are central to *Zeitgeist*'s operation. Our goal is to exploit only the topology of cross-references, rather than the raw text of the corresponding Wikipedia articles (which would necessitate heavy-duty parsing and analysis methods). Since some topological contexts are more opaque than others, *Zeitgeist* employs a multi-pass approach to acquiring new word forms. In the first pass, only clear-cut cases are harvested; these exemplars are then generalized to underpin schemata that, in a second pass, allow less obvious neologisms to be recognized and semantically analyzed. Both passes are described in sections three and four. In section five, an empirical evaluation and discussion of *Zeitgeist*'s results is presented, while concluding thoughts are offered in section six.

2 LINKING WORDNET AND WIKIPEDIA

WordNet and Wikipedia each blur the traditional semiotic distinction between dictionaries and encyclopedias – which views the former as a source of *word* knowledge and the latter as a source of *world* knowledge – in different ways. WordNet is primarily an electronic dictionary/thesaurus whose structure is informed by psycholinguistic research (e.g., it uses different representations for nouns, verbs, adjectives and adverbs), but in eschewing alphabetic indexing for a semantic organization, it imposes an encyclopedia-like topic organization on its contents. Its coverage is broad, containing entries on topics such as historical events, places and personages more typically found in an encyclopedia. Unsurprisingly, it tends to be used in NLP applications not just as a lexicon, but as a lightweight knowledge-base for reasoning about entities and events.

For its part, Wikipedia's topic articles are surprisingly word-oriented. One finds many more headwords than in a conventional encyclopedia, and a richer level of interconnectedness. In many cases, composite headwords (such as "feminazi") are explicitly linked to the entries for their component parts, while detailed articles on lexical phenomena such as blended (or portmanteau) word-forms [4, 5] and political epithets provide links to numerous topical

[1] School of Computer Science and Informatics, University College Dublin, Belfield, Dublin, Ireland.

examples. Additionally, a sister project, Wiktionary [6], aims to exploit the Wikipedia model for an open-source dictionary.

The advantages accruing from an integration of such complementary resources are obvious. To Wikipedia, WordNet can give its explicit semantic backbone, as found in the *isa*-taxonomy used to structure its noun senses. To WordNet, Wikipedia can give its rich, open-textured topology of cross-references [7], as well as its larger and constantly growing set of topical headwords. To achieve this integration, the headwords of Wikipedia must be sense-disambiguated, though [8] report positive results for this task. In this paper, we explore the extent to which the semantic head of a neologism (that part which contributes the suffix, partially or completely, such as "pub" in "Gastropub" and "economics" in "Enronomics") can be disambiguated by the priming effects of other links emanating from the same Wikipedia article. General purpose WSD techniques (e.g., [9,10]), applied to the text rather than the links of an article, can then be used to resolve those ambiguous heads that are not primed in this way.

Toward thus end, we introduce two connectives for relating Wikipedia headwords to WordNet lexical entries. The first is written x *isa* y, and states that a new synset {x} is to be added to WordNet as a hyponym of the appropriate sense of y. Thus, *superhero isa hero* assumes that WSD is used to identify the intended sense of "hero" in the "superhero" context. The second is x *hedges* y, as in *spintronics hedges electronics*. As described in Lakoff [11], a hedge is a category-building relationship that allows one to reason as if a concept belonged to a given category, in spite of strict knowledge to the contrary (e.g., most people know that whales are not fish, but reason about them as if they were). In WordNet terms, hedge relationships will ultimately be instantiated via taxonomic coordination: {spintronics} will not be added as a hyponym of {electronics}, rather both will share the common hypernym {physics}. Hedges allow us to sidestep the awkward issues of hyperbolae and metaphor that frequently mark new coinages. Though "affluenza" ("affluence + influenza") is not, strictly speaking, a kind of "influenza", the hedge allows an NLP system to reason as if it were a real virus; this is apt, since the blend is used to depict affluence as a contagious affliction.

3 PASS I: LEARNING FROM EASY CASES

We employ a string-matching approach to recognizing and analyzing Wikipedia neologisms, in which specific schemata relate the form of a headword to the form of the words that are cross-referenced in the corresponding article. Let $\alpha\beta$ represent the general form of a Wikipedia term, where α and β denote arbitrary prefix and suffix strings that may, or may not, turn out to be actual morphemes. In addition, we use $\alpha \to \beta$ to denote a reference to headword β from the Wikipedia article of α, and use $\alpha \to \beta\,;\gamma$ to denote a contiguous pair of references to β and γ from article α.

As noted earlier, Zeitgeist seeks out neologisms that are a formal blend of two different lexical inputs [4, 5]. The first input contributes a prefix element, while the second contributes a suffix element that is taken to indicate the semantic head of the neologism as a whole.

The first schema below illustrates the most common arrangement of lexical inputs (as we shall see in section 5):

Schema I: Explicit Extension

$$\frac{\alpha\beta \to \beta \;\wedge\; \alpha\beta \to \alpha\gamma}{\alpha\beta \textit{ isa } \beta}$$

This schema recognizes blended word forms like "gastropub" and "feminazi" in which the suffix β is a complete word in itself (e.g., "pub" and "Nazi"), and in which the prefix α is a fragment of a contextually linked term (like "gastronomy" or "feminist"). The suffix β provides the semantic head of the expansion, allowing the new term to be indexed in WordNet under the appropriate synset (e.g., {Nazi} or {pub, public_house}). The textual gloss given to this new entry will be a simple unpacking of the blended word: "$\alpha\gamma$ β" (e.g., "gastronomy pub" and "feminist Nazi"). To avoid degenerate cases, α and β must meet a minimum size requirement (at least 3 characters apiece), though in some exceptional contexts (to be described later), this threshold may be lowered.

Many neologisms are simple variations on existing terminology. Thus, "fangirl" is a male variation on "fanboy", while "supervillain" is a criminal variation on "superhero". When an explicit Wikipedia reference exists between these alternating suffixes, the new composite word can be identified as follows:

Schema II: Suffix Alternation

$$\frac{\alpha\beta \to \alpha\gamma \;\wedge\; \beta \to \gamma}{\alpha\beta \textit{ hedges } \alpha\gamma}$$

This schema identifies a range of alternating suffix pairs in Wikipedia, from *man*↔*boy* to *woman*↔*girl* to *genus*↔*genera, bit*↔*byte* and *bacteria*↔*toxin*.

We can now begin to consider portmanteau words in which the suffix term is only partially present. Words like "Rubbergate" are understood as variations on other terms (e.g., "Watergate") if the prefix term (here, "rubber") is explicitly linked. In effect, a partial suffix like "gate" becomes evocative of the whole, as follows:

Schema III: Partial Suffix

$$\frac{\alpha\beta \to \gamma\beta \;\wedge\; (\alpha\beta \to \alpha \vee \alpha\beta \to \delta \to \alpha)}{\alpha\beta \textit{ hedges } \gamma\beta}$$

This schema additionally covers situations where the prefix is only indirectly accessible from the neologism, as in the case of "metrosexual" (where "metro" is accessible via a link to "metropolitan"), and "pomosexual" (where "pomo" is only accessible via a mediating link to "postmodernism"). We note that this schema ignores the obvious role of rhyme in the coinage of these neologisms.

This indirection means that, in words like "metrosexual", both the prefix and the suffix may be partially projected to form a true portmanteau word. In Wikipedia, the lexical inputs to a portmanteau word are often stated as contiguous references in the corresponding article. For instance, Wikipedia describes "sharpedo" as a "shark torpedo", while "Spanglish" is explicitly unpacked in the corresponding article as "Spanish English". We can exploit this finding in the following schema:

Schema IV: Consecutive Blends

$$\frac{\alpha\beta \rightarrow \alpha\gamma\ ;\ \delta\beta}{\alpha\beta\ \textit{hedges}\ \delta\beta} \qquad \textit{e.g., sharpedo} \rightarrow \textit{shark torpedo}$$

Indeed, portmanteau terms are so striking that the corresponding Wikipedia articles often explicitly reference the headword "portmanteau", or vice versa. In such cases, where $\alpha\beta \rightarrow$ *portmanteau*, we can safely reduce the minimum size requirements on α and β to two characters apiece. This allows *Zeitgeist* to analyze words like "spork" (spoon + fork) and "sporgery" (spam + forgery).

4 PASS II: RESOLVING OPAQUE CASES

The foregoing schemata anchor themselves to the local topological context of a headword to curb the wild over-generation that would arise from string decomposition alone. But even when this topological context is uninformative, or absent entirely (since some Wikipedia articles make no reference to other articles), a system may be able to reason by example from other, more clear-cut cases. For instance, there will be many exemplars arising from schemas III and IV to suggest that a word ending in "ware" is a kind of software and that a word ending in "lish" or "glish" is a kind of English. If E is the set of headwords analyzed using schema III and IV, and S is the corresponding set of partial suffixes, we can exploit these exemplars thus:

Schema V: Suffix Completion

$$\frac{\alpha\beta \rightarrow \gamma\beta \wedge\ \gamma\beta \in E\ \wedge\ \beta \in S}{\alpha\beta\ \textit{hedges}\ \gamma\beta}$$

Since the Wikipedia entries for "crippleware", "donationware" and "malware" – but not "stemware" or "drinkware" – make reference to "software", the above schema allows us to infer that the former are kinds of software and the latter dishware. Suffix completion reflects the way neologisms are often coined as reactions to other neologisms; for example, once "metrosexual" is recognized using schema III (partial suffix), it provides a basis for later recognizing "retrosexual" using schema V, since "sexual" will now suggest "metrosexual" as a completion. Similarly, "Reaganomics" serves as an exemplar for later analysing "Enronomics".

If P denotes the set of prefix morphemes that are identified via the application of schemas I, II and III, we can also formulate the following generalization:

Schema VI: Separable Suffix

$$\frac{\alpha\beta \rightarrow \beta \wedge \alpha \in P}{\alpha\beta\ \textit{isa}\ \beta} \qquad \textit{e.g., antiprism} \rightarrow \textit{prism}$$

This is simply a weakened version of schema I, where α is recognized as a valid prefix but is not anchored to any term in the topological context of the headword.

Though the entry "logicnazi" makes no reference to other headwords in Wikipedia, one can immediately recognize it as similar to "feminazi" (a "feminist Nazi" as resolved by schema I). Conceptually, "Nazi" appears an allowable epithet for an extreme believer of any ideology, and in part, this intuition can be captured by noting that the "Nazi" suffix overwrites the "ism" / "ist" suffix of its modifier. If T is a set of tuples, such as <ism, Nazi>, derived from the use of schema I, we have:

Schema VII: Prefix Completion

$$\frac{\alpha\gamma \rightarrow \alpha\ \wedge <\gamma, \delta\beta> \in T}{\alpha\beta\ \textit{isa}\ \beta}$$

Zeitgeist recognizes "logicnazi" as a kind of "Nazi", in the vein of "feminazi", since, from "logic" it can reach an "ism" or belief system "logicism" for this Nazi to extol. Likewise, it recognizes "Zionazi" as an extreme Zionist (allowing for a shared "n"), and "Islamonazi" as an extreme Islamist (allowing for an added "o" connective).

Finally, the collected prefixes and suffixes of pass one can now be used to recognize portmanteau words that are not explicitly tagged (as in schema V) or whose lexical inputs are not contiguously referenced (as in schema IV):

Schema VIII: Recombination

$$\frac{\alpha\beta \rightarrow \alpha\gamma\ \wedge \alpha\beta \rightarrow \delta\beta \wedge \alpha \in P \wedge \beta \in S}{\alpha\beta\ \textit{hedges}\ \delta\beta}$$

Thus, a "geonym" can be analyzed as a combination of "geography" and "toponym".

5 EVALUATION AND DISCUSSION

To evaluate these schemata, each was applied to the set of 152,060 single-term headwords and their inter-article connections in Wikipedia (as downloaded as a SQL loader file in June, 2005). Version 1.6 of WordNet was used to separate known headwords from possible neologisms. In all, 4677 headwords are decomposed by one or more of the given schemata; of these: 1385 (30%) are ignored because the headword already exists in WordNet, 884 (19%) are ignored because the hypernym or hedge determined by the analysis does not itself denote a WordNet term. Thus, though "bioprospecting" is correctly analyzed as "biology

prospecting", "prospecting" is not a lexical entry in WN1.6 and so this term must be ignored. The remaining 2408 (51%) of cases[2] are analyzed according to the breakdown of Table I:

Table 1. Breakdown of performance by individual schema.

Schema	# Headwords		# Errors	Precision
I	710	29%	11	.985
II	144	5%	0	1.0
III	330	13%	5	.985
IV	82	3%	2	.975
V	161	6%	0	1.0
VI	321	13%	16	.95
VII	340	14%	32	.90
VIII	320	13%	11	.965

Each *Zeitgeist* analysis was manually checked to find errors of decomposition and provide the precision scores of Table I. Two schemas (II in pass one, which e.g., derives Rubbergate from Watergate, and V in pass two, which e.g., derives retrosexual from metrosexual) produce no errors, while the most productive schema (explicit extension, schema I) has an error rate of just 1.5%. In contrast, schema VII (prefix completion in pass two, which derives logicnazi via the exemplar feminist/feminazi) is cause for concern with an error rate of 10%. High-risk schemata like this should thus be used in a controlled manner: they should not update the lexicon without user approval, but may be used to hypothesize interpretations in contexts that are more ephemeral and where more information may be available (e.g., a spellchecking or thesaurus application invoked within a particular document).

Some obvious factors contribute to an overall error rate of 4%. Company names (like Lucasfilm) comprise 12% of the erroneous cases, organization names (like Greenpeace and Aerosmith) 6%, place names (like Darfur) 11% and product names (like Winamp) 2%. Another 5% are names from fantasy literature (like Saruman and Octopussy). In all then, 35% of errors might be filtered in advance via the use of a reliable named-entity recognizer.

5.1 Word Sense Disambiguation

For 51% of the Wikipedia neologisms recognized by *Zeitgeist*, the semantic head (i.e., the word that contributes the suffix to the neologism) denotes an unambiguous WordNet term. The remaining 49% of cases thus require some form of WSD to determine the appropriate sense, or senses, of the semantic head before the neologism can be added to WordNet. While one can employ general purpose WSD techniques on the textual content of a Wikipedia article [9, 10], the topological context of the headword in Wikipedia may, to a certain degree, be self-disambiguating via a system of mutual priming.

For example, the intended WordNet sense of "hero" in the headword "superhero" (not present in WN 1.6) is suggested by the link *superhero*→ *Hercules*, since both "hero" and "Hercules" have senses that share the immediate WordNet hypernym {Mythological-Character}. In general, a given sense of the semantic head will be primed by any Wikipedia term linked to the neologism that has a WordNet sense to which the head relates via synonymy, hyponymy or hypernymy.

Priming can also be effected via an intersection of the textual glosses of WordNet senses and the topological context of the Wikipedia article (in a simple Wikipedia variation of the Lesk algorithm [9]). For example, the Wikipedia headword "kickboxing" suggests the ambiguous "boxing" as a semantic head (via schema I). However, because the Wikipedia link *kickboxing*→*fist* is echoed in the gloss of the WordNet sense {boxing, pugilism, fisticuffs} but not in the gloss of {boxing, packing}, only the former is taken as the intended sense.

More generally, the elements of the Wikipedia topological context can be viewed as a simple system of semantic features, in which e.g., *fist* is a feature of *kickboxing*, *fascism* is a feature of *Nazi*, and so on. Furthermore, because blending theory [4,5] claims that blended structures will contain a selective projection of elements from multiple inputs, this projection can be seen in the sharing of semantic features (that is, topological links) between the neological headword and its semantic head. For instance, the Wikipedia terms "Feminazi" and it semantic head, "Nazi", share three Wikipedia links – to Totalitarianism, Fascism and Nazism – which may be taken as the contribution of the lexical component "Nazi" to the meaning of the word as a whole. In the terminology of blending theory [4,5], these features are *projected* from the input space of Nazi into the blended space of Feminazi. Projection of this kind occurs in 64% of the neologisms recognized by *Zeitgeist*.

By understanding the projective basis of a word blend, *Zeitgeist* has yet another means of performing disambiguation of the semantic head, since the intended sense of the head will be that sense that visibly contributes semantic features to the blend. In the case of "kickboxing", the feature *fist* is directly contributed by the pugilistic sense of "boxing". However, for the blended word "emoticon", the feature *pictogram* is indirectly contributed by the user-interface sense of "icon" via its hypernym {symbol}.

Overall, topological priming resolves 25% of neologisms to a single WN1.6 sense, while another 1% are resolved to multiple WN senses, which is to be expected when the head element is a polysemous word. For instance, "photophone" ("photograph" + "telephone") is deemed to hedge both the equipment and medium senses of "telephone", while "subvertising" ("subversion" + "advertising") is deemed to

[2] Interestingly, the distribution for WN2.1 is much the same: 1570 analysed headwords (33%) are ignored because the headword is already in WN2.1, while 789 headwords (17%) must be ignored because their semantic heads are *not* in WN2.1. This leaves 2319 valid neologisms (49%) to be added to WN2.1, as opposed to 2408 for WN1.6. The number of neologisms remains relatively stable across WN versions because greater lexical coverage presents a greater opportunity to recognize neologisms that cannot be integrated into lesser versions. For instance, the "cyberpunk" entry in WN2.1 means that while this word is not treated as a neologism for this version (as it is for WN1.6), its presence allows "steampunk" and "clockpunk" to be recognized as neologisms.

hedge the message and industry senses of "advertising". In all, total WSD coverage in Zeitgeist is 77%. Recourse to more general WSD techniques is thus needed for just 23% of cases.

5.2 Literal Versus Figurative Interpretations

Our evaluation reveals that over half (57%) of the neologisms recognized by *Zeitgeist* (via schemas I, VI and VII) are realized in WordNet via a simple hypernymy relationship, while the remainder (43%) are realized (via schemas II, III, IV, V and VII) using the more nuanced *hedge* relationship. It seems clear, for instance, that "Gastropub" really is a kind of "pub" and "cocawine" really is a kind of "wine" (with added cocaine). However, it is not so clear whether Feminazis are truly Nazis (in the strict, National Socialist sense), so hedging may be more prevalent than these figures suggest. Though WordNet defines {Nazi} as a hyponym of {fascist}, the word is often used as a highly charged pseudo-synonym of the latter. "Nazi" seems to be used here in a sense-extensive, metaphoric fashion to suggest totalitarian zeal rather than political affiliation.

Two factors alert us that this use of "Nazi" is hyperbolae rather than literal extension. The first is the orthographic form of the word itself, for while "Nazi" is a proper-named class, "Feminazi" employs the word in an uncapitalized form which suggests a process of semantic bleaching or generalization. The second factor is the relative contribution, in terms of projected features, of the semantic head to the blend as a whole. Recall that the word "Nazi" shares the Wikipedia linkages {Totalitarianism, Fascism, Nazism} with "Feminazi", so these features may be said to originate from this input. However, "fascist" also references the terms {Totalitarianism, Fascism, Nazism} in Wikipedia, suggesting that there is no obvious loss of semantic import if Feminazi is considered an extension of {fascist} rather than of {Nazi}.

In 36% percent of neologisms, one or more semantic features are projected into the blend by a hypernym of the semantic head. In just 2% of neologisms this projection occurs in the context of an *isa* relation (i.e., via schemas I and VI) and is such that all features that are projected from the head are also redundantly projected from the hypernym of the head. (As it happens, only in the case of "Feminazi" does the semantic head denote a proper-named concept). While not conclusive, such redundancy is sufficient cause either to hedge the relationship or to prompt for human guidance in these cases.

6 CONCLUSIONS

We have presented a linguistics-lite approach to harvesting neologisms from Wikipedia and adding them to WordNet. *Zeitgeist* does not employ an explicit morphological analyser, but relies instead on a marriage of partial string-matching and topological constraints. Nonetheless, many of the words that are successfully recognized exhibit a creative and playful use of English morphology. Furthermore, by grounding is analyses in the local link topology of Wikipedia articles, *Zeitgeist* gains a semantic insight that one cannot obtain from morphology rules alone. For instance, not only is "microsurgery" recognized as a micro-variant of surgery, the specific meaning of "micro" in this context is localized to the headword "microscopy" via schema I. The concept "microsurgery" is not just "micro-surgery", but surgery conducted via a microscope.

Even a lightweight approach can, however, bring some degree of semantic insight to bear on the analysis of new words. In this respect, Wikipedia's link topology deserves further consideration as a source of semantic features. Certainly, Wikipedia has great promise as a semi-structured semantic representation. For instance, one can distinguish two kinds of semantic feature in Wikipedia. Strong or highly-salient features are those that are reciprocated; thus, *charity*→*altruism* and *altruism*→*charity* implies that altruism is a highly salient feature of charity, and vice versa. Weak features are those that are not reciprocated in this way. It remains to be seen how far one can go with such a representation without imposing a more rigid logical framework, but we believe that the initial foray described here suggests the scheme has yet more mileage to offer.

We conclude by noting that the linguistics-lite nature of *Zeitgeist's* approach means that is not intrinsically biased toward English. In principle, its mix of string matching and topological constraints should validly apply to other languages also. Whether phenomena like lexical blending spring forth with equal regularity in the non-English languages supported by Wikipedia is a subject of future research.

REFERENCES

[1] Dent, S., Fanboys and Overdogs: The Language Report III. Oxford University Press (2003).

[2] Miller, G. A., WordNet: A Lexical Database for English. *Communications of the ACM*, Vol. 38 No. 11 (1995).

[3] *www.wikipedia.org*

[4] Fauconnier, G., Turner, M., Conceptual Integration Networks. *Cognitive Science*, 22(2), pp 133–187 (1998).

[5] Veale, T., O'Donoghue, D., Computation and Blending. *Cognitive Linguistics*, 11(3-4).

[6] *www.wiktionary.org*

[7] Ruiz-Casado, M., Alfonseca, E., Castells, P., Automatic Extraction of Semantic Relationships for WordNet by Means of Pattern Learning from Wikipedia. *LNAI 3513, pp 67 (2005).*

[8] Ruiz-Casado, M., Alfonseca, E., Castells, P., Automatic Assignment of Wikipedia Encyclopedic Entries to WordNet Synsets. Springer *LNAI 3528, pp 280 (2005).*

[9] M. Lesk, Automatic sense disambiguation using machine readable dictionaries: how to tell a pine cone from an ice cream cone. *Proceedings of ACM SigDoc Conference, Toronto: ACM*, pp 24-6 (1986).

[10] P. Resnik, Semantic Similarity in a Taxonomy: An Information-Based Measure and its Application to Problems of Ambiguity in Natural Language. *Journal of Artificial Intelligence Research* 11, pp 95-130, (1999).

[11] G. Lakoff, *Women, Fire and Dangerous Things: How the Mind forms Categories*. University of Chicago Press (1987)