ARTICLE IN PRESS

Available online at www.sciencedirect.com

SCIENCE DIRECT®

Knowledge-Based Systems xxx (2006) xxx–xxx

Knowledge-Based SYSTEMS

www.elsevier.com/locate/knosys

An analogy-oriented type hierarchy for linguistic creativity

Tony Veale

School of Computer Science and Informatics, University College Dublin, Belfield, Dublin 4, Ireland

Received 23 January 2006; accepted 15 April 2006

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. We demonstrate 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.

Keywords: Analogy; Ontologies; Type hierarchies; WordNet; Creativity

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 [6]. 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 [6,16] and the words we use to communicate these categories [13,14]. This fluidity contrasts sharply with the rigidity of the taxonomies that have been traditionally posited to organize our category systems [8,11,23]. Our goal in this paper is to show how fluidity can be considered an emergent property of existing taxonomies when types are treated as dynamic, rather than static, entities that are created as they are needed.

Taxonomies have, since antiquity [8], 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* [8], 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

E-mail address: Tony.Veale@UCD.ie.

0950-7051/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.knosys.2006.04.007

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 exemplified by [23].

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 [23] 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 [23] gives as an example "Nixon is the submarine of world politics", and suggests that this metaphor is resolved by the dynamic type *Things Which Behave In A Secret Or Hidden Manner*. 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 unresolved key issues of exhaustiveness and computational tractability.

Now, analogies [15] and metaphors [6,13] exhibit creativity not just in their production – since they constitute novel linguistic artifacts – but also in their interpretation, which frequently requires the dynamic construction of new ad-hoc categories (of the form championed by Barsalou [2]). Consider the rather cryptic joke offered by Freud [5], and analyzed explicitly as an analogy in [1]:

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

The analysis in [1] provides the missing concept, Prostitute, to complete the analogy:

(1) *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 an old-fashioned Freudian view of marriage and wives), as well as Personal-Resource and Hired-Resource. The burden of creativity is not borne alone 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.

2. Dynamic types as ad-hoc categories

The Freudian analogy/joke of (1) clearly demonstrates that type creation is a crucial sub-process of not just analogy construction, but of analogy interpretation also. From this perspective, the goal of a good analogy is to prompt the listener to dynamically create a P-Creative type (see Refs. [3,4]) that already exists (perhaps as a H-Creative innovation) within the lexical ontology of the speaker. In other words, many dynamic types are ad-hoc by-products [2] of the analogical reasoning mechanism. Consider the analogies in (2):

(2a) *Zeus:Greek :: Jupiter:Roman*
(2b) *Zeus:Greek :: ???:Roman*
(2c) *"Zeus is the Greek Jupiter"*

The analogy of (2a) 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 (2b) employs an elliptical form common to tests 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 (2c) assumes a compressed natural language form that can also be considered a metaphoric expression [8].

The implicit relation common to (2a), (2b) and (2c) appears to "deity of": Zeus is a deity of the Greeks, while Jupiter is a deity of the Romans. However, consider a longer form of this analogy in (3):

(3) *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 (3) 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: first, the concept must already exist, yet is not lexicalized; or two, neither the concept nor its lexicalization exist 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 (4), which can be seen as a close conceptual neighbor of (2a):

(4) *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 the whose lexicalization "war god" is such an conventional collocation that few would argue that it is constructed especially for the analogy. However, this is not to say that the interpretation of (4) should be substantially different from that of (2) or (3). 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 (4), this relational category (War-God) will correspond to an existing lexical concept, while in (2) and (3) it will result in a lexical innovation ("supreme deity") that may or may not be added to the lexicon following either 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 (e.g., Ref. [14,17]) we should adopt. To remain as agnostic as possible, let us assume a rather simple, feature-theoretic view of the lexicon. Let $\boldsymbol{F}$ 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 $\uparrow$, 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:

$$\boldsymbol{F}(Zeus) = \{\uparrow deity,\ @Greek, @supreme, @mythology, @Olympus\}$$

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

$$\boldsymbol{F}(Jupiter) = \{\uparrow deity,\ @Roman, @supreme, @mythology, @rain\}$$

Jupiter is thus a deity that is Roman and supreme, associated with both mythology and rain (in the guise of *Jupiter Pluvius*).

Since we are dealing with a lexico-conceptual representation of words and concepts, any new type must possess both a lexical and a conceptual realization. Since a new type can be arbitrarily complex in either of these respects (recall our discussion of Way [23] and her unrestrained dynamic types), we make the simplifying assumption (to be reconsidered in Section 5) that any new type is lexically realized as a compound term comprising a taxonomic head and a domain modifier. The set of allowable modifiers is given by precisely those terms that have been used as modifiers in existing lexical concepts (and which can thus be applied to other heads). Thus we introduce the following operator:

$\mathbf{U_M}\{\mathbf{X}\}$: *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, ...\}

Likewise, the set of allowable heads is given by precisely those terms that have been used as heads in existing lexical concepts (and which can thus be meaningfully differentiated):

$\mathbf{U_H}\{\mathbf{X}\}$: *Usage as head*: return a set of all compound terms 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, ...\}

Any dynamically created type will therefore combine a modifier term (with proven differentiation ability) with a head term (with proven scope for differentiation) as follows:

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

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

The feature-level decomposition of lexico-conceptual structure, given by $\boldsymbol{F}$, thus 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 $\boldsymbol{F}$:

$$\boldsymbol{alt}(\mathrm{A}) = \boldsymbol{C}(\{\mathrm{X} \mid \mathbf{U_M}\{\mathrm{X}\} \neq \{\} \wedge @\mathrm{X} \in \boldsymbol{F}(\mathrm{A}) \wedge \uparrow \mathrm{X} \notin \boldsymbol{F}(\mathrm{A})\}, \{\mathrm{Y} \mid \mathbf{U_H}\{\mathrm{Y}\} \neq \{\} \wedge \uparrow \mathrm{Y} \in \boldsymbol{F}(\mathrm{A})\})$$

That is, the set of alternate categorizations of A comprises just those compound terms that can be created by combining the associative features of A that have in the past been used as 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 the respect to the lexicon and the language that it represents.

A new type represents an ad-hoc category for a given task if it is created in the context of that task in the furtherance of that task [2]. The set of ad-hoc categories arising from a given analogy A:B, which we denote ***adhoc***(A:B), is thus a simple function of the set of alternate categorizations of A and B, as follows::

$$\begin{aligned}&\boldsymbol{adhoc}(\mathrm{A} : \mathrm{B})\\ &\quad = \{\mathrm{X} - \mathrm{Y} \mid \mathrm{X} - \mathrm{Y} \in \boldsymbol{alt}(\mathrm{A}) \cap \boldsymbol{alt}(\mathrm{B}) \wedge \\ &\qquad \neg(\exists \mathrm{P} \uparrow \mathrm{P} \in \boldsymbol{F}(\mathrm{A}) \wedge \uparrow \mathrm{P} \in \boldsymbol{F}(\mathrm{B}) \wedge \uparrow \mathrm{Y} \in \boldsymbol{F}(\mathrm{P}))\}\end{aligned}$$

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 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:

$$\begin{aligned}&\boldsymbol{wellformed}(\mathrm{A}:\mathrm{B}::\mathrm{C}:\mathrm{D})\\&\quad=(\exists \mathrm{M}\,\mathrm{M}-\mathrm{B}\in \boldsymbol{alt}(\mathrm{A})\wedge \mathrm{M}-\mathrm{D}\in \boldsymbol{alt}(\mathrm{C}))\\&\qquad\vee(\exists \mathrm{H}\,\mathrm{B}-\mathrm{H}\in alt(A)\wedge D-H\in \boldsymbol{alt}(\mathrm{C}))\\&\qquad\vee(\exists \mathrm{M}_1\mathrm{M}_2\mathrm{H}_1\mathrm{H}_2\,\mathrm{M}_1-\mathrm{H}_1\in \boldsymbol{alt}(\mathrm{A})\\&\qquad\wedge \mathrm{M}_1-\mathrm{H}_2\in \boldsymbol{alt}(\mathrm{C})\wedge \mathrm{M}_2-\mathrm{H}_1\in \boldsymbol{alt}(\mathrm{B})\\&\qquad\wedge \mathrm{M}_2-\mathrm{H}_2\in \boldsymbol{alt}(\mathrm{B}))\end{aligned}$$

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:

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

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 (2a) and (4), it is sufficient that this be a non-empty set. But in the case of long-form analogies like (3), 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 (2b), 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:

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

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.

3. 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 Refs. [20,21]): 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 Hebrew 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): = {$\uparrow$Greek-deity $\uparrow$deity @supreme @god @ancient @mythology}
F(Jupiter): = {$\uparrow$Roman-deity $\uparrow$deity @Roman @ supreme @god @Romans}
F(Alpha): = {$\uparrow$letter $\uparrow$character @1st @letter @Greek @alphabet}
F(Aleph): = {$\uparrow$letter $\uparrow$character @1st @letter @ Hebrew @alphabet}

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

adhoc(Zeus:Jupiter) = {supreme-deity}
adhoc(Alpha:Aleph) = {1st-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, 1st-letter is deemed the stronger of the two. In fact, an extensional analysis reveals that the extension of 1st-letter (with just two members) is a proper subset of that of alphabet-letter (with 49 members), which suggests that 1st-letter is a specialization of the category alphabet-letter.

Fig. 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.

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

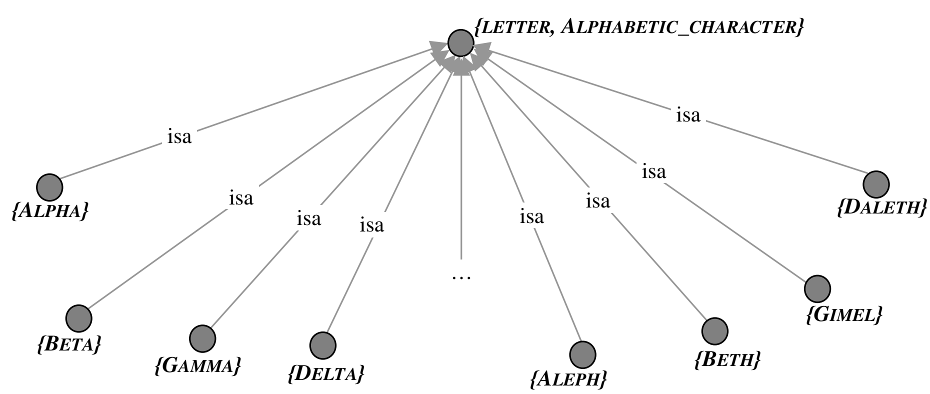

Fig. 1. The structure of the Greek- and Hebrew-letters domain in WordNet.

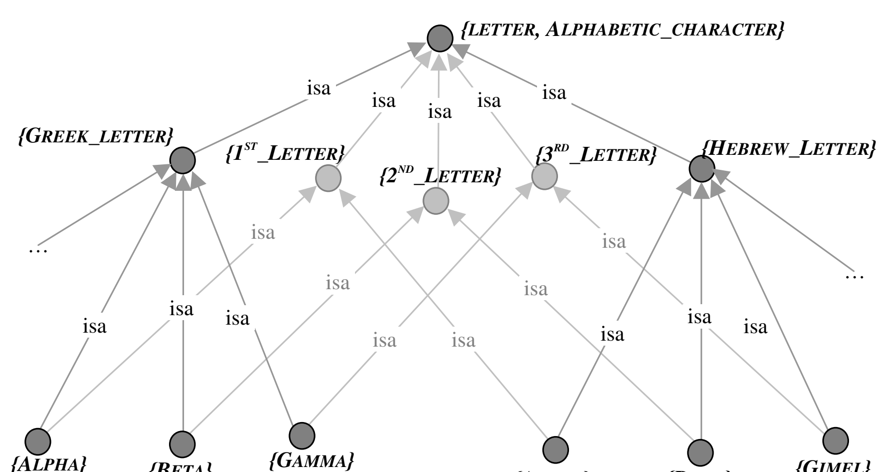

Fig. 2. WordNet supplemented with new adhoc categories like Greek-letter, Hebrew-letter and 1st-letter, created as by-products of analogical retrieval.

Table 1
Analogical Mappings between different deity pantheons as represented in WordNet (version 1.6)

Adhoc basis	*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	Ler
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

as the latter lacks vowels, a strict 1-to-1 alignment is not possible. Fig. 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.

4. 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.

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).

4.1. 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:

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

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. 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 an average of three different ways.

Overall, the most differentiating modifier is ‘‘Mexico’’, which serves to differentiate 34 different head; for example, 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. 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.

5. Palimpsest: A new ontological overlay for WordNet

We should now reconsider the simplifying assumption that makes possible the creation of these new taxonomic types. For in stipulating that a new type will be lexicalized as a two-word combination of a domain modifier and a taxonomic head, we seriously limit the form of the generalizations that can be induced (e.g., see Refs. [10,12]), allowing these generalized types to make use of only the simplest atomic features, like *supreme*, 1st and *Hebrew*. In many cases an atomic feature is sufficiently discriminating to support an analogy between two lexical concepts (as demonstrated by the experiments in Section 4), but in many others, inter-concept similarity is predicated on shared properties that are considerably more complex. Consider the lexical concepts Work-Crew and Jury: both are social groups (and represented in WordNet as such), but their analogical similarity extends further than this. In English, the leader of a work-crew is called a "foreman", as is the leader of a jury, which suggests that the concept Jury is commonly conceived of as a group of workers (a "work detail") that is assembled for a specific purpose and then disbanded. To capture this analogical similarity, a system would need to induce a new type of the form *Groups-whose-leaders-are-called-Foreman*, which seems to take us perilously close to the baroque typing of Way's [23] DTH. At the very least, the discrimination underlying this shared property seems beyond the simple view of the lexicon (via ***F***) offered here, and most likely beyond the information content of the glosses offered by WordNet.

Features like *supreme*, 1st and *Hebrew* are essentially adjectival in nature, but in many cases the necessary discriminating property for an analogy will a multi-place relation that is lexicalized as a verb. For instance, the WordNet lexical concepts Deliveryman, Roundsman, Deliverer and Bailor (none of which are synonymous) are similar by virtue of sharing the relational property *delivers*:*merchandise*. Such relational types can in turn be generalized further, to support an analogy between Paperboy and Milkman (each deliver something, suggesting the abstracted type *delivers*:?), between Cobbler and Plumber (each repair something, suggesting the abstracted type *repairs*:?), and so on. These relations might be extracted from WordNet by parsing the corresponding glosses (e.g., see Refs. [7,18]) but WordNet glosses are neither regular enough in form, or consistent enough in content, to make this approach a practical reality.

These issues lead us to a conclusion that is at once both dispiriting and liberating: the semantic basis for these new types cannot reliably be extracted via automatic means, but must be added manually. Certainly, such an effort is time-consuming, but ultimately, we believe it is less so than the automated approach, which is in constant need of tweaking to overcome problems of under-generation and over-generalization [17,18]. To ensure a comprehensive coverage, and to allay fears that this effort is a representational *toy* designed to work only for specially selected examples, we conceive of the effort as a new lexical ontology in its own right, but one that is to be overlaid on the existing taxonomic structure of WordNet. The new ontology, named Palimpsest, will attempt to offer an explicit property-theoretic description of each lexical concept in WordNet, in a form that will allow for the dynamic generation of new types as they are needed.

Each Palimpsest property is associated with a specific WordNet synset, and relates that synset/concept to another WordNet lexical entry. Properties comprise both a predicate/relation and an object, and may be marked in a variety of ways to indicate salience. The general form of a property is given as follows:

$$[*]\text{relation}[/\text{abstraction}] : \text{object}[*]$$

Each property can be seen as specifying a category of concepts that all share that property. Thus, Jury and Crew both share the property *leader*/*manager*:*Foreman*, which becomes the de facto name for the ontological type that embraces both concepts. Within types, each relation may be specified relative to a common abstraction that is shared by multiple relations. Thus, the relations *cleans*/*restores* and *repairs*/*stores* can be seen as specializations of the abstract relation restores, while the relations *sells*/*provides* and *pays*/*provides* are specializations of the abstract relation *provides*. This two-level specification of relations allows a system to perform functional abstraction on a concept, to e.g., recognize the similarity between cleaners, repairmen and surgeons (see Ref. [22]).

The optional * at either end of the property mark either the relation (if before) or the object (if after) as especially salient and foregrounded within the definition of the concept. For instance, in the definition of Merchant, the relation *sells*/*provides*, and the object, merchandise, are both foregrounded, as follows:

$$\{\text{merchant, merchandiser}\} : *\text{sells/provides} \\ : \text{merchandise}*$$

When an object is foregrounded in this way, it means that it (e.g., Merchandise) is central to the definition of the host concept. When a relation is foregrounded, it means that the host concept (e.g., Merchant) can be used in place of that relation to convey that relation in another (possibly analogical or metaphoric) context. For instance, another concept in which the relation *sells*/*provides* is foregrounded is Peddler; we can thus reconceptualize the concept Merchant as a "merchandise peddler", a Stockbroker as a "stock merchant" or "stock peddler", a prostitute as a "sex merchant", and so on.

Dynamic types arise in Palimpsest when properties are generalized or composed. For instance, the property *sells*/*provides*:*merchandise* can be generalized either as *provides*:*merchandise*, *sells*/*provides*:? or *sells*/*provides*:?. Whenever a property is generalized, its corresponding category/type is widened to admit new members, permitting those new members to be seen as similar in ways that were not previously perceived. In contrast, composition is type-narrowing operation, allowing multiple properties to be conjoined into one, as when two properties in the same host concept share the same object, allowing the corresponding relations to be merged. For instance, a Dealer buys and sells merchandise, while a Stockbroker buys and sells shares. In each case, a new composite relation can be created and then generalized, to yield *buy*_and_ it sell:?. Properties can be composed in a variety of other ways (e.g., hierarchically), allowing new types with more exclusive category memberships to be dynamically constructed on the fly.

Palimpsest is still at an experimental and immature stage of development. Only with maturity will it settle on an appropriate set of relations and relational abstractions, and maturity will only come via scale. However, development is progressing apace, and Palimpsest, now nearing the 10,000 concept mark, will soon be ready for general release to the research community (where it will be free for research purposes).

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. Like the most striking and novel of metaphors, a good analogy can reveal deep insights between the most far-removed concepts, like wives and umbrellas or prostitutes and taxi-cabs [1,4]. To be able to generate such insightful categories is a grand and ambitious goal in its own right within the rubric of creativity research, but such insights, if achieved through automated means, can also alleviate the structural problems that inevitably plague manually constructed ontologies. This is especially true of those ontologies 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 have exploited, 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 lexical 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.

Nonetheless, there are fundamental limits to the automation of dynamic behaviour in a lexico-conceptual system like WordNet, due in large part to the intrinsic representational limits of WordNet itself (as variously discussed in Ref. [17–21]). Though WordNet glosses are a rich source of implicit propositional structures, they lack sufficient regularity and completeness to fuel a full solution to the problem of new type creation. We acknowledge then the need to supplement these automatic mining techniques with a large-scale hand-coding effort (much like that which yielded WordNet in the first place) that will provide the explicit structures from which new types can arise. This effort, dubbed Palimpsest, is currently underway and will soon reach the 10,000 concept mark, at which point it will be released to the research community at large.

References

[1] S. Attardo, C.F. Hempelmann, S. Di Maio, Script oppositions and logical mechanisms: modeling incongruities and their resolutions, Humor, Int. J. Humor Res. 15 (1) (2002) 3–46.
[2] L. Barsalou, Ad-Hoc Categories, Mem. Cogn. 11 (3) (1983) 211–227.
[3] M. Boden, The Creative Mind, Abacus, London, 1990.
[4] M. Boden, Creativity: a framework for research, Behav. Brain Sci. 17 (3) (1994) 558–568.
[5] S. Freud, Der Witz und seine Beziehung zum Unbewußten. (Dueticke, Leipzig 1905; Reprinted Frankfurt am Main: Fischer. English edition,1976. Jokes and their Relation to the Unconscious. Harmondsworth: Penguin.).
[6] G. Lakoff, M. Johnson, Metaphors we live by, University of Chicago Press, 1980.
[7] S.M. Harabagiu, G. A. Miller and D. I. Moldovan, WordNet 2 – A Morphologically and Semantically Enhanced Resource, in: Proceedings of the ACL SIGLEX Workshop on Standardizing Lexical Resources (Maryland, 1999).
[8] J. Hutton, Aristotle's Poetics, Norton, New York, 1982.
[9] D. Lenat, R.V. Guha, Building Large Knowledge-Based Systems, Addison Wesley, Reading Massachusetts, 1990.
[10] J. McCarthy, Creative Solutions to Problems, in: Proceedings of the AISB'99 symposium on AI and Scientific Creativity (Edinburgh Scotland, 1999).
[11] G.A. Miller, WordNet: a lexical database for English, Commun. ACM 38 (11) (1995).
[12] S. Muggleton, Inductive Logic Programming, New Generat. Comput. 13 (1991) 245–286.
[13] A. Ortony, The role of similarity in similes and metaphors, in: A. Ortony (Ed.), Metaphor and Thought, Cambridge University Press, Cambridge, UK, 1979.
[14] J. Pustejovsky, The generative lexicon, Comput. Linguist. 17 (4) (1991).
[15] T. Veale, M.T. Keane, The competence of sub-optimal structure mapping on hard analogiesProceeding of the IJCAI'97, the International Joint Conference on Artificial Intelligence, Nagoya, Japan, Morgan Kaufman, San Mateo California, 1997.
[16] T. Veale, D. O'Donoghue, Computation and Blending, Cognitive Linguistics 11 (3/4) (2000) 253–281.
[17] T. Veale, Systematicity and the Lexicon in Creative Metaphor, in: Proceedings of the ACL 2003 workshop on Figurative Language, (Sapporo, Japan, 2003).
[18] T. Veale, Qualia Extraction and Creative Metaphor in WordNet, in: Proceeding of the 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 CommunicationLNAI 2746, Springer, Berlin, 2003.

[20] T. Veale, The analogical thesaurus: an emerging application at the juncture of lexical metaphor and information retrieval, in: Proceedings of the IAAI'03, The AAAI Press, Menlo Park California, 2003.
[21] T. Veale, Creative Information Retrieval, in: Proceedings of the CICLing'04, Springer Verlag LNCS, Berlin, 2004.
[22] T. Veale, Analogy Generation in HowNet, in: Proceedings of the IJCAI'05, Morgan Kaufman, San Mateo California, 2005.
[23] E.C. Way, Knowledge Representation and Metaphor, Studies in Cognitive systems, Kluwer Academic Publishers, Amsterdam, 1991.