Dynamic Type Creation in Metaphor Interpretation and Analogical Reasoning: A Case-Study with WordNet

TonyVeale

Department of Computer Science,
University College Dublin, Belfield, Dublin 6, Ireland.
Tony.veale@UCD.ie
http://www.cs.ucd.ie/staff/tveale/home/

Abstract. Metaphor and Analogy are perhaps the most challenging phenomena 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. This research considers the problem of how a conceptual system structured around a central taxonomy can dynamically create new categories or types to understand novel metaphors and analogies. This theoretical perspective yields a practical method for dynamic type creation within WordNet.

1 Introduction

This paper considers the dynamic nature of the representations required to facilitate two of the most knowledge-hungry processes in a conceptual system, metaphor interpretation and analogical reasoning. These processes are interesting because they often exploit the latent similarities between domains that have not been explicitly represented in the underlying conceptual structure [8], thus revealing the inadequacies of such structures. In particular, because metaphors and analogies are used to create new ways of thinking about familiar things, they reveal the essential fluidity of the categories we use to structure the world [9]. This fluidity contrasts sharply with the rigidity of the taxonomies that have been traditionally employed to organize our category systems [4].

Taxonomies have, since antiquity [1], provided a systematic means of hierarchical decomposition of knowledge, whereby a domain is successively dissected via differentiation into smaller pockets of related concepts. Rich differentiation leads to effective clustering, so that similar concepts become localized to the same region of the taxonomy. This locality not only makes the categorial similarity of different ideas 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* [1], 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 [2], a common-sense ontology for general reasoning, and WordNet [3], a psycholinguistically motivated lexical database of English. The Aristotelian view of taxonomic metaphor also continues to exert considerable influence in computer theories, as demonstrated by [4] and [5].

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 as a tool for affecting change in our category systems. To resolve this contradiction, authors such as Eileen Cornell Way [4] have argued for the importance of a dynamic type hierarchy (DTH) as a taxonomic backbone for conceptual structure. Such a taxonomy would dynamically reveal new types in response to appropriate metaphors. For example, Way [4] 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, Way does not suggest an empirical means of constructing such a DTH, which effectively leaves the issue of exhaustiveness, and all it entails for computational tractability, unresolved.

This paper describes an empirical means of constructing a DTH that dynamically generates new taxonomic types in response to challenging analogies and metaphors. The underlying static type hierarchy is provided by WordNet 1.6, while dynamic types are extracted when needed from the textual glosses provided by the designers of WordNet[1]. In addition, we identify an important class of taxonomic type we dub an "analogical pivot", and show how types in existing taxonomies like WordNet and Cyc, which contain relatively few such pivots naturally, can be automatically upgraded into pivots by the addition of dynamic types, further facilitating the processes of analogical retrieval and mapping.

---

[1] Note that WordNet does not enforce a type/instance distinction, and so uses the *isa* relation to denote both subset relations and membership relations.

2 Analogical Pivots

Taxonomic systematicity implies that related or 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. For example, in WordNet 1.6, the concept {*alphabet*}[2] is differentiated culturally into {*Greek_alphabet*} and {*Hebrew_alphabet*}, but the concept {*letter, alphabetic_character*} is not similarly differentiated into {*Greek_letter*} and {*Hebrew_letter*}. Rather, 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*). But more than similarity judgment is impaired: crucially, a lack of systematicity and symmetry in differentiation undermines another core rationale of taxonomic structure, the ability to recognize analogies and metaphors. For instance, the structure of WordNet 1.6 is not sufficiently differentiated for analogies like "what is the Jewish gamma?" *(gimel)* or who is the "Viking Ares" *(Tyr)*.

Consider the analogical compound "Hindu Zeus" and how one might interpret it using 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*}. Now, because WordNet also defines an entry for {*Hindu_deity*}, it requires just a simple composition of ideas to determine that the "Hindu Zeus" will be daughter of the type {*Hindu_deity*}. More generally, we simply find the lowest parent of the head term ("Zeus") that, when concatenated with the modifier term ("Hindu") or some synonym thereof, yields an existing WordNet concept. We dub this type, here {*deity, god*}, the pivot of the analogy, since the mapping process can use this pivot to construct a target counterpart of the source concept that significantly narrows the space of possible correspondences. So the Hindu counterpart of Zeus is a daughter of {*Hindu_deity*}, and the precise one can be chosen on the basis of other types (such as {*supreme_deity*}) that encode finer differentiating criteria. Fig. 1 presents a schematic view of this process.

---

[2] Each concept in WordNet is unambiguously denoted by a *synset* of synonymous words that can all be used to denote the same underlying concept. We do not list all the members of a synset when it clear what entity/type is being denoted.

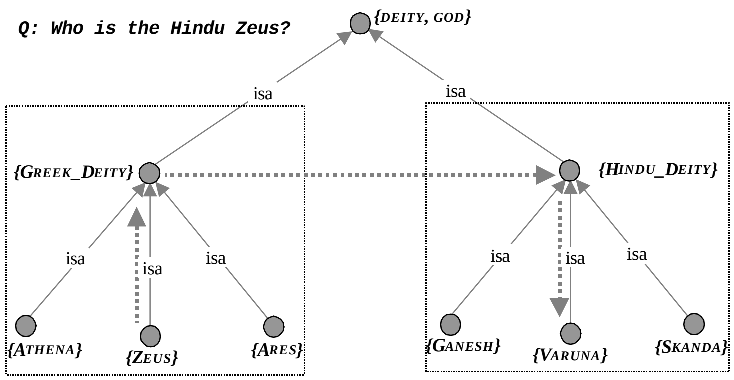

Fig. 1. An example of taxonomic structure driving an analogy between domains

Compare this approach with the conventional one of taxonomic reconciliation, due to Aristotle[1], in which two entities or types can be considered analogous if they share a common superordinate. This approach still finds considerable traction in computational models today (e.g., see [4,5]), but it is easily trivialized: in a well designed taxonomy, any two entities or types will always share at least one superordinate (even if it is the root type), and so any two concepts will always be potential analogues in such a system. The current approach uses a much stricter notion of taxonomic analogy: two types are potentially analogous if they each possess superordinates that are themselves analogous differentiations of the same direct parent (the pivot of the analogy). The approach is also constructive: it indicates how the target counterpart of the pivot, {*Hindu_deity*} is to be constructed from the source analogue (Zeus). Thus, Zeus and Varuna are analogous because {*Greek_deity*} and {*Hindu_deity*} are analogous, by virtue of being different domain sub-types of the same pivot. This constraint is the taxonomic equivalent of the squaring rule described in [6] to ensure that there is structural support for every analogical mapping.

2.1 Conceptual Association

More formally, an analogical pivot is an interior type of a taxonomy whose daughter types (hyponyms) are explicitly differentiated to extend into different domains of knowledge. Pivots thus sit at the junctures of different domains and act as the most effective signposts into those domains when performing an analogical mapping. So the key to the mapping process is to first locate the pivot of the analogy and then follow the labeled signposts into the target domain. For robust understanding, every synonym of the target modifier should be considered when following these signposts, so that “Hindoo Zeus” or even

"Hindustani Zeus" will point the same target taxonomy. However, for the mapping to effectively navigate from a pivot in this way, we require a more powerful notion than simple synonymy.

We thus broaden the notion of synonymy to that of symmetric associativity. By definition, synonyms are symmetric associates of each other since one can be substituted for the other without loss of meaning (e.g., "Moslem" and "Muslim"). We broaden this notion to include terms that are so closely correlated in meaning that one can be used as a metonymic proxy for the other (e.g., "Muslim", "Islam", and "Koran"). This associativity can be determined statistically from a corpus, but a simpler and more principled method involves using the WordNet entries themselves.

Definition: The *symmetric associates* of a word X comprises the set of synonyms of X, as well as the set of each word Y that appears in a definition/gloss of a sense of X such that X also appears in the definition of an individual sense of Y (**1**)

Thus "Islam" is a symmetric associate of "Muslim" since the former occurs in a definition of the latter and vice versa. Similarly by this reckoning, the symmetric associates of "Hindu" are {*Hindu, Hindoo, Hinduism, Hindustan Hindustani, Trimurti*}, where "Trimurti" denotes a triad of divinities in Hindu mythology. By using the symmetric associations of the target to differentiate the pivot, the mapping process can understand analogies as allusive as "Trimurti Zeus".

The above approach is still very fragile as natural pivots like {*deity, god*} are extremely rare in WordNet, since it has not been explicitly constructed for analogical purposes. As noted earlier, the WordNet concept {*letter, alphabetic_character*} is not culturally differentiated, so a mapping cannot be constructed for "Jewish delta" → {*Hebrew_letter*}. Furthermore, even when pivots do exist to facilitate the mapping, what is produced is a target hypernym rather than a specific domain counterpart. One still needs to go from {*Hindu_deity*} to {*Varuna*} (like Zeus a supreme cosmic deity, but of Hinduism), or from {*Hebrew_letter*} to {*daleth*} (like "delta" the fourth letter, but of the Hebrew alphabet).

3 Creating Dynamic Types

Both problems can be solved by adding additional differentiating structure to the taxonomy. These new types will dynamically dissect the taxonomy in novel ways as new metaphors and analogies attempt to establish points of comparison between domains. The dynamic creation of these types will, in the process, convert the parents of these types into the analogical pivots that are needed to direct the interpretation of analogies and metaphors from the source to the target domain.

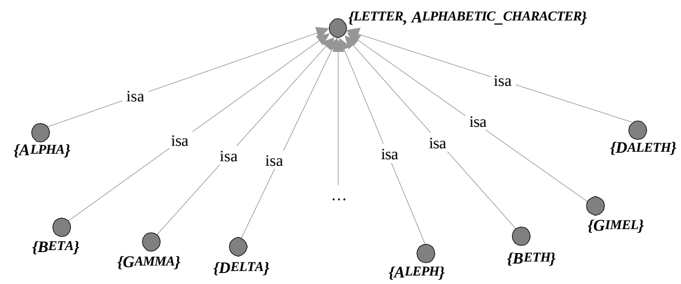

Fig. 2. The impoverished *{letter, alphabetic_character}* taxonomy in WordNet 1.6

For example, the creation of new types like {*Greek_letter*} and {*Hebrew_letter*} will transform {*letter, alphabetic_character*} into a pivot that extends into the Greek and Hebrew domains. Types such as these act as signposts from the pivot into specialized areas of the taxonomy and thus allow the first cut of the analogical mapping to occur. In contrast, other dynamic types may be less ambitious: a new type like {*1st_letter*} will unite just two concepts, {*alpha*} and {*aleph*}, and a new type like {*thunder_deity*} will also unite just two concepts, {*thor*} and {*donar*}. However, these lower-level types allow for finer-grained mapping within the target domain once the appropriate area of the taxonomy has been identified using the pivot.

Fig. 2 illustrates the situation as found in WordNet 1.6 as it concerns the representation of the {*letter, alphabetic_character*} domain. Note how no differentiating structure exists within the taxonomy to allow an analogizer to even discriminate the letters of one alphabet from the other, which makes mapping of domain counterparts impossible. Compare this structure with that of Fig. 3, which illustrates the most desirable state of the taxonomy from an analogy and metaphor perspective. In this reworking, the {*letter, alphabetic_character*} concept is not only differentiated by domain into {*Greek_letter*} and {*Hebrew_letter*}, it is simultaneously differentiated by relative letter position. This structure is sufficient to allow a near-isomorphic mapping to be generated from one alphabet to another (with the exception of extra letters 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.

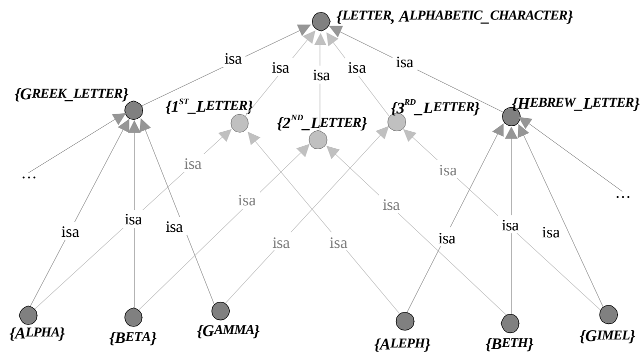

Fig. 3. The taxonomic structure of *{letter, alphabetic_character}* becomes a richly structured lattice when enriched with a variety of new types like *{Greek_letter}* and *{1^{st}_letter}*

We do not refer to new types like {*1^{st}_letter*}, or even {*Greek_letter*}, as pivots; rather, it is their existing parent, {*letter, alphabetic_character*}, that becomes a pivot after these types have been added. Once a type has been sufficiently differentiated by a number of sub-types, it can act as a sign-posted crossroads between multiple domains, and thus facilitate precise mappings of entities from those domains. When enough pivots are in place, a taxonomy becomes a decision lattice for metaphor and analogy. In Fig. 3 above, only one entity lies at the intersection of {*Hebrew_letter*} and {*1^{st}_letter*}, only one at the intersection of {*Hebrew_letter*} and {*2^{nd}_letter*}, and so on, so with this lattice, a 1-to-1 mapping of entities from one alphabet to another can be generated. It is also possible to convert such a lattice into a decision tree by labeling the arcs of the taxonomy appropriately. One such decision tree for a fragment of the {*deity, god*} sub-taxonomy in WordNet 1.6 is illustrated in Fig. 4.

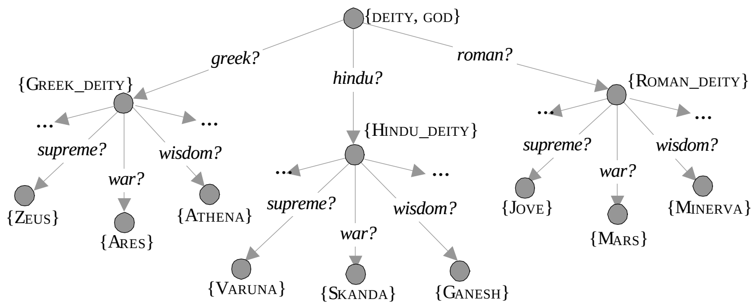

Fig. 4. A decision-tree perspective on the sub-taxonomy for *{deity, god}* in WordNet

3.1 Type Creation As Feature Reification

Enhancing the differentiating power of WordNet is essentially a task of feature reification. WordNet (like other taxonomies, such as Cyc [2]) expresses some of its structure explicitly, via isa-links, and some of it implicitly, in textual glosses intended for human rather than machine consumption. Fortunately, these glosses are consistent enough to permit automatic extraction of structural features (e.g., see [7], who extract lateral connections between concepts from these glosses). What is needed is a means to recognize the word features in these glosses with the most analogical potential, so that they may be lifted to create new taxonomic types. The noun sense glosses of WordNet 1.6 collectively contain over 40,000 unique content words, but clearly only a small fraction of these words can be profitably reified. We thus employ two broad criteria to identify the words worth reifying, *differentiation potential* and *alignment potential*:

Definition: A lemmatized word-form has *differentiation potential* if it occurs in more than one gloss, but not in too many glosses (e.g., more than 1000). Additionally, there must be a precedent for using the word as an explicit differentiator in at least one existing taxonomic entry (2)

Definition: A word-form has *alignment potential* if it can be found in multiple places in the taxonomy at the same relative depth from a pivot (3)

Consider the word "wisdom", which occurs in 20 different WordNet glosses, enough to demonstrate cross-domain potential but not too many to suggest over-generalization. Additionally, there is a WordNet precedent, {*wisdom_tooth*}, for its explicit use as a differentiator. And of the concepts that "wisdom" is used to gloss, at least three – {*Athena*}, {*Ganesh*} and {*Minerva*} – are grand-daughters of the concept {*deity, god*}. As shown in Fig. 5., the word "wisdom" thus has alignment potential relative to the concept {*deity, god*}, suggesting that "wisdom" can be reified to create a new taxonomic concept {*wisdom_deity*}.

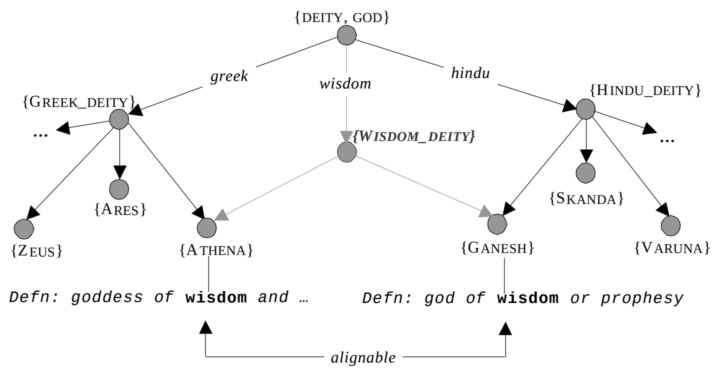

Fig. 5. Analysis of the gloss for *{Athene}* suggests that the word-form "wisdom" has analogical potential, since it is alignable with another use in *{Ganesh}*

How does one identify the potential pivot against which alignability is measured? In general, any interior type of the taxonomy can be a potential pivot, but from a practical perspective, it makes sense to only consider the atomic types that have not already been differentiated. Thus, {*deity, god*} is a potential pivot but {*Greek_deity*} is not, since the latter is already domain specific. We thus assume the following:

Definition: A hypernym X is a potential pivot relative to a hyponym Y if X is the lowest, undifferentiated (atomic) hypernym of Y (4)

Thus, when we consider the word forms in the gloss of {*Athena*}, alignability will be determined relative to the concept {*deity, god*} rather than {*greek_deity*}, so that any reification that is performed will create a new differentiation of the former. With an appropriate reverse-index of gloss words to the concepts that are defined by them, this makes the identification of alignable features very efficient. The system simply needs to examine each concept reachable via the index entry for the word and consider only those at the same relative depth from the potential pivot.

4 Roles and Relations

Analogies that are interpreted in taxonomic terms have the advantage of being conceptually justified via the process of type inclusion, since the source and target concepts are understood relative to the same common super-type. There is thus a sound conceptual basis for considering the target to be an analogue of the source. Furthermore, the extent of this basis can be quantified numerically by considering how far one must ascend in the taxonomy to find this common super-type.

Yet this advantage is bought at the price of symmetry: because taxonomic interpretations are created on the basis of commonalities, they tend to be highly symmetric, while the most creative interpretations tend to be highly asymmetric [8]. Insightful metaphors and analogies help to inform and enrich the target concept by imposing the highly-developed relational structure of the source onto the less-developed target [9]. Without this imposition of relational structure, metaphor is reduced to the status of a fanciful but redundant way of referring to existing concepts, which, in WordNet at least, already have linguistic labels. Systematic projection of structure from one domain to another is the basis of the structure-mapping approach to analogy and metaphor [6], which determines correspondences between entities based on their relative positions in a larger relational structures.

It thus becomes necessary to determine the relational structure of each concept if a full metaphoric/analogical interpretation is to be generated. Of course, one can build these representations by hand, but tedium aside, the resulting structures would have little theoretical force, since any conceptual theory can be made to appear tractable if one has complete freedom to hand-craft its representations. For this reason, we prefer instead to extract such representations automatically from independent linguistic data. There are two broad sources of data for this task: one can look to external corpora and attempt to mine relational patterns from large quantities of raw text, or one can look to the lexicographer glosses that accompany most concepts in WordNet. In either case, some type of parsing must be done on the linguistic data to extract relational patterns, and the reliability of extraction will depend crucially on the robustness of the parsing method used. We choose to mine the glosses, since their direct association with specific concepts resolves some of the problems of lexical ambiguity that can arise.

We describe here an extraction technique that balances coverage with quality: by attempting to extract a relatively narrow slice of the relational structure inherent in WordNet glosses, we can be confident of quite high levels of competence. We thus concentrate on relations that pertain to the agency/telicity of a concept, both because they can be extracted reliably and because they tend to capture the most intrinsic, behavioral aspects of a concept that are most likely to be projected by a metaphor.

4.1 Extraction of Relational Structure

The extraction technique is simple to define and straightforward to apply. Simply, it combines a knowledge of the derivational morphology of English with the taxonomic structure of WordNet, so that extracted relations are both linguistically and conceptually sound. The key here is that the agent-telic aspects of a concept are often expressed using nominalized verbs that implicitly encode relational structure, such as "observer" (from "observe") and "specializer" (from "specialize"), and these nominalization patterns can be captured in a small number of highly productive morphology rules. For example, concepts such as {*geologist*} and {*linguist*} are defined with glosses that explicitly invoke the term "specializer", while {*witness*} is defined relative to the term {*observer*}. Now, WordNet 1.6 provides two senses for "specializer", of types {*doctor*} and {*expert*}, both of which are sub-types of {*person*}. The concepts {*geologist*} and {*linguist*} are also sub-types of {*person*}, strongly suggesting that "specializer" is an appropriate telic relation for each. Note how the WordNet taxonomy plays a key role in recognizing this relation. If "specializer" did not have a sense that was compatible with {*person*}, it would be rejected as a relation, so ultimately, extraction depends crucially on the taxonomic metaphor.

It is also straightforward to morphologically derive the patient form of "specializer" as "specialism", allowing a system to conclude that both {*geologist*} and {*linguist*} have the relational structure *specializer_of:specialism* (we leave the particular sense of "specializer" under-specified for maximal metaphoric reuse). Now, while strongly telic nouns like "specializer" are often used in WordNet glosses, the underlying verbs themselves that are even more frequent. For example, the concepts {*surgeon*} and {*pastry_cook*} are both provided with glosses that use the word "specializes", but using the same morphology rules in reverse, the corresponding nominalization "specializer" can be found. In this way both concepts are receive the relational structure *specializer_of:specialism*.

Using morphological rules in conjunction with taxonomic type checking, a large quantity of agent-telic relations can be robustly extracted from glosses with the simplest of shallow parses. Broad clues to the syntactic form of the gloss (such as active versus passive voice) can be derived from a combination of keyword analysis and inflectional morphology. The passive voice causes a relational arc to be inverted, as in the case of {*dupe*}, whose gloss is *"a person who is swindled or tricked"*. The resulting relational structure is thus: *of_swindler:swindler* ∧ *of_trickster:trickster*.

The glosses of many WordNet concepts suggest a metonymic relational structure. Consider the gloss assigned to the concept {*diary, journal*}: *"a daily record of (usually private) experiences and observations"*. The morphology of the word "diary" itself yields the agentive relation *of_diarist:diarist*, while nominalization rules suggest the additional relations *of_experience:experience*, *recorder_of:recording* and *observer_of:observation*. However, doubt is cast upon the latter two by subsequent taxonomic analysis, which reveals that {*diary, journal*}, a sub-type of {*communication*}, is not compatible with either {*recorder*} or {*observer*}, both sub-types of {*person*}. Nonetheless, these relations are not rejected, since the suggested patients, {*recording*} and {*observation*}, like {*diary, journal*}, are sub-types of {*communication*}, suggesting that a diary can be seen as a kind of metonym for the observer/recorder (as evoked by the familiar address "dear diary"). The concept {*diary, journal*} therefore yields the relational structure *meta_observer_of:observation* ∧ *meta_recorder_of:recording* ∧ *of_experience:experience*.

4.2 Projection of Relational Structure

The projection of relational structure can be performed either literally or figuratively. In a literal interpretation, the relational structure of the source is simply instantiated with the target concept, so for example, a literal "travel diary" is a diary that contains travel recordings and travel observations. In contrast, figurative interpretations first attempt to find a target domain correspondence for the source concept, and then project the relational structure of the source onto this counterpart [6]. For instance, WordNet suggests {*passport*} as a figurative reference for "travel diary" since both are kinds of travel document. Projecting the relational structure of {*diary, journal*} onto {*passport*} causes the latter to be seen as a journal of travel observations and experiences, and indeed, many travelers retain old passports for this very purpose. In structure-mapping terms, the more productive the metaphor, the greater the amount of relational structure that is projected. Furthermore, the more systematic and apt the metaphor, the greater the isomorphism between the structure that is projected from the source onto the target [6, 9].

4.3 Aptness and Projection

Metaphors are most apt when projection highlights a latent relational structure that already exists in the target concept [8]. For example, the compound "pastry surgeon" can be understood taxonomically as referring to {*pastry_cook*}, since like {*surgeon*} it is a sub-type of {*person*}. While this seems quite appropriate, the taxonomic approach arrives at the precisely the same interpretation even when the compound is "pastry astrologer" or "pastry hostage". To see why {*surgeon*} is more apt than {*astrologer*} as a source concept, one need look no further than their relational structures. WordNet 1.6 defines a surgeon as a *"physician who specializes in surgery"*, while a pastry cook is glossed as *"a chef who specializes in pastry"*. Both {*surgeon*} and {*pastry_cook*} contain the relation *specializer_of:specialism*, and it is precisely this relation that is highlighted by the metaphor. In contrast, the concept {*astrologer*} causes the relations *astrologer_of:astrology, forecaster_of:forecast* and *visionary_of:vision* to be projected, and neither of these relations exist in the {*pastry_cook*} structure.

Metaphors also appear more apt when they systematically evoke, or connect into, established modes of metaphoric thought [9]. Consider the compound "political mechanic": many different concepts can be reached from "political" that prove to be taxonomically compatible with the concept {*mechanic*}, among them {*political_leader*}, {*political_scientist*} and {*machine_politician*}. However, the extracted structure of {*mechanic*} contains the relation *machinist_of:machine*, whose surface similarity with {*machine_politician*} is highly suggestive. More interestingly, however, the instantiated structure for "political mechanic" thus becomes *machinist_of:political_machine*, where {*political_machine*} is a conventional metaphor already established within WordNet. This marks "political mechanic" as a systematic outgrowth of an established metaphor, making it seem all the more appropriate. Comparable systematicity is exhibited by the compounds "political chemist", which relationally connects to {*political_science*}, "political missionary", which in WordNet 1.6 connects to {*political_program*}, and "political torchbearer", which connects to {*political_campaign*}.

5 Empirical Analysis

Dynamic types are created in the context of specific metaphor interpretation or analogical reasoning tasks. For example, types like {*Hebrew_letter*} and {*Greek_letter*} are created in response to specific analogies, such as "What is the Jewish delta?". However, to test the applicability of the type creation process, we have pre-applied this type creation process to 69,780 unique noun senses in

WordNet 1.6, whose glosses collectively contain 35,397 unique unlemmatized content words.

Now, because of the strict reification criteria for feature-lifting from glosses (see definitions 2 and 3), only 2806 of these content words are reified, to add 9822 new types, like {*cheese_dish*}, to WordNet. These types serve to differentiate 2737 existing concept in WordNet, such as {*dish*}, transforming these concepts into analogically-useful pivots. In total, 18922 noun concepts (27% of the sample) are connected to the new types, via the addition of 28,998 new isa-links to WordNet. Each dynamic type thus serves to unite an average of 3 daughters apiece. But in a subsequent pass over the new types, 1258 additions (or 12.8%) were culled because they did not sufficiently differentiate their parents to be worthwhile. For example, the type {*Greek_gorgon*} is worthless since all known gorgons are Greek.

A review of the other 87.2% of differentiators reveals that WordNet is being dissected in new and useful ways, both from the perspective of simple similarity judgments (e.g., the new types achieve a fine-grained clustering of similar ideas) and from the perspective of analogical potential. Overall, the most differentiating feature is "Mexico", which serves to differentiate 34 different pivots (such as *{dish}*, to group together {*taco*}, {*burrito*} and {*refried_beans*}), while the most differentiated pivot is {*herb, herbaceous_plant*}, which is differentiated into 134 sub-categories (like {*prickly_herb*}). To consider just a few other domains: sports are differentiated into team sports, ball sports, court sports, racket sports and net sports; constellations are divided according to northern and southern hemispheric locales; 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 further differentiated to yield {*war_deity*}, {*love_deity*}, {*wine_deity*}, {*sea_deity*}, {*thunder_deity*}, {*fertility_deity*}, and so on.

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 is unavailable for cultural reasons, N/A is used to fill the corresponding cell. In two cases, marked by (*), an adequate mapping could not be generated when one was culturally available; 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 solely defined 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 letters of the Greek alphabet are mapped onto the 23 letters of the Hebrew alphabet (as represented in WordNet), and vice versa. The recall rate for the Hebrew to Greek letter task, for both dynamic and static WordNet hierarchies, is 100%, while for the reverse task, Greek to Hebrew, it is 96% (since Greek contains an extra letter). However, the precision of the static hierarchy is only 4%, since every letter of the target alphabet appears equally similar as a candidate mapping (Fig. 2), while for the dynamic hierarchy it is 96% (Greek to Hebrew) and 100% (vice versa).

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

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

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 (e.g., the Norse to Hindu mapping 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 defined as “supreme cosmic deity” in WordNet 1.6, with no explicit reference of Hinduism). Average precision for the dynamic hierarchy is 93.5%, with the loss of 6.5% precision due to 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.

6 Conclusions

Manually constructed representations on the ambitious scale of WordNet and Cyc are naturally prone to problems of incompleteness and imbalance. The one-size-fits-all nature of the task results in a taxonomy that is often to undifferentiated for precise similarity judgments and too lopsided to support metaphor and analogical mapping. A symptom of this incompleteness is the fact that English glosses or commentaries provide the ultimate level of differentiation, so that one cannot truly differentiate two concepts without first understanding what the glosses mean. The goal of this work is to lift the implicit discriminators out of the flat text of the glosses and insert them into the taxonomy proper, as dynamic types that will facilitate finer similarity judgments and richer analogical mappings.

It is interesting to note that WordNet 1.6 already contains lexical entries for some dynamic types, such as {*war_god*} and {*sea_god*}, but fails to relate them to the appropriate sub-types such as {*Ares*} and {*Poseidon*}. This may be due to simple oversight, but is more likely a symptom of WordNet's tendency toward single-inheritance (even though multi-inheritance is supported, and actually used in a relatively small number of noun concepts). With single inheritance, the taxonomist must decide between differentiating a concept on the basis of culture (e.g., Greek) or on that of function (e.g., war), rather than simply doing both. Whatever the true cause, dynamic types do more than facilitate metaphor and analogy, but actually repair deficiencies in an existing static taxonomy.

References

1. Hutton, J.: Aristotle's Poetics. Norton, New York (1982)
2. Lenat, D., Guha, R.V.: Building Large Knowledge-Based Systems. Addison Wesley (1990)
3. Miller, G. A.: WordNet: A Lexical Database for English. Communications of the ACM, Vol. 38 No. 11 (1995)
4. Way, E. C.: Knowledge Representation and Metaphor. Studies in Cognitive systems, Kluwer Academic Publishers (1991)
5. Fass, D: An Account of Coherence, Semantic Relations, Metonymy, and Lexical Ambiguity Resolution. In: Small, S. I, Cottrell, G. W., Tanenhaus, M.K. (eds.): Lexical Ambiguity Resolution: Perspectives from Psycholinguistics, Neuropsychology and Artificial Intelligence. Morgan Kaufman, San Mateo California (1988)
6. Veale, T., Keane, M. T.: The Competence of Sub-Optimal Structure Mapping on Hard Analogies. The proceedings of IJCAI'97, the Int. Joint

Conference on Artificial Intelligence, Nagoya, Japan. Morgan Kaufman, San Mateo California (1997)
7. Harabagiu, S. M., Miller, G. A., Moldovan, D. I.: WordNet 2 - A Morphologically and Semantically Enhanced Resource. The Proceedings of the ACL SIGLEX Workshop: Standardizing Lexical Resources. Maryland, USA (1999)
8. Ortony, A.: The role of similarity in similes and metaphors. In: Ortony, A. (ed.): Metaphor and Thought. Cambridge University Press : Cambridge, U.K (1979)
9. Lakoff, G., Johnson, M.: Metaphors we live by. Uni. of Chicago Press: Chicago (1980)
10. Pustejovsky, J.: The generative lexicon. Computational Linguistics, Vol. 17 No. 4 (1991)