

# Understanding Metaphor Using a Corpus-trained Fluid Category System

Yanfen Hao and Tony Veale<sup>1</sup>

<sup>1</sup> School of Computer Science and Informatics, University College Dublin, Ireland  
{yanfen.hao, tony.veale}@ucd.ie

**Abstract.** Metaphor and analogy are processes that creatively play with our conventional conceptions of what constitutes a category structure. Knowledge of conventional category boundaries can be found in ontologies of varying richness and coverage, from WordNet to Cyc. But what each of these ontologies have in common is a rather static, topdown view of category structure. In this paper we outline a dynamic, bottom-up view of category membership based on context-sensitive corpus analysis. By learning from corpora about how people creatively use categories, we believe that our computational systems can learn for themselves how to replicate this category-level creativity.

## 1 Introduction

Metaphor has long been viewed as an ontological means of creatively shifting our conception of a topic from one category to another, from a category to which it is assigned by convention to one where certain desirable properties can appear more salient [5,8,10]. The key question in metaphor research is whether this ontological shift is a superficial sleight of hand played out at the level of words, or a conceptual operation that directly affects the ontological representation of ideas. Proponents of the former assume that metaphors are rhetorical devices that can always be substituted with a more literal paraphrase of the speaker's intended meaning [9]. Proponents of the latter instead argue that metaphors are precisely what they appear to be – statements about category membership – and should be interpreted as such, to directly affect to how these categories are cognitively structured [3,8]. Thus, the metaphoric statement “my job is a jail” is not a cue to see my job as merely similar to a jail, but a category-broadening proposition that views a job as an actual member of the category Jail, albeit a member that is highly atypical and far from prototypical [5].

This latter view of [3] requires an extremely flexible, non-classical view of category structure, one that views category membership as a graded rather than binary notion [5], one in which concepts can fluidly move from one category to another [4]. But this fluidity does not sit well with conventional perspectives on ontological structure, as represented by the ontologies of [2,6,7]. In this paper we look at one conventional ontology, the HowNet system of [2], which is a large-scale bilingual lexical ontology for words and their meanings in both Chinese and English. In many

respects, HowNet is similar to the WordNet lexical ontology for English [7], though in contrast to WordNet, HowNet provides an explicit, if sparse, propositional semantics for each of the word-concepts it defines. Complementing this frame-like semantics, in which concepts are defined in terms of actions, case-roles and fillers, is a taxonomic backbone that seems rather impoverished when compared to that of WordNet. Nonetheless, we argue that to effectively process analogies and metaphors, a new taxonomic structure must be derived from HowNet’s propositional semantics. Since this derivation process is automated, it allows us to re-invent HowNet in non-classical terms, so that the categories that comprise its new taxonomic backbone have graded membership criteria [5] and fluid boundaries [4] that can admit new members in figurative contexts [3]. In effect, we describe here how HowNet can be reformulated as a SlipNet in the style of Hofstadter’s work with fluid analogies [4].

In section two, we describe how an analogy-oriented taxonomy of concepts, based on how entities actually function and behave, can be derived from HowNet. As described in [10], this new taxonomy still maintains the classical structure categories with binary membership criteria and sharp boundaries. In section three, however, we describe how this classical structure can be made more fluid, to become a Hofstadter-style SlipNet in which concepts can slide from one category to another. In section four we augment this structure with a sense of directionality, since metaphor is an asymmetric phenomenon in which category shifts are highly directional [8]. Section five then provides a preliminary empirical evaluation of this work, before some concluding remarks are offered in section six.

## 2 Analogical Reasoning With HowNet

HowNet and WordNet each reflect a different view of semantic organization. WordNet [7] 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, or synsets, and further differentiates these synsets from one another by assigning them to different positions of a taxonomy. In contrast, HowNet is constructive in nature. It does not provide a human-oriented textual gloss for each lexical concept, but instead composes sememes from a less discriminating taxonomy to provide a semantic representation for each word sense. For example, HowNet defines the lexical concept surgeon|医生 as follows:

(1) 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 an agent of a doctoring activity” (the {~} here serves to indicate the placement of the concept within its associated propositional structure). We see a similar structure employed by HowNet for the lexical concept repairman|修理工:

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

Note that the impoverished nature of HowNet’s taxonomy means that over 3000 different concepts are forced to share the immediate hypernym human|人. However, human|人 merely states, very generally, what a repairman is, rather than what a repairman does. Fortunately, HowNet also organizes its verb entries taxonomically, and so we find the verbs doctor|医治 and repair|修理 organized under the hypernym resume|恢复 (the logic being, one supposes, that “doctoring” and “repairing” both involve a resumption of an earlier, better state). This similarity of verbs, combined with an identity of case-roles (both surgeon and repairman are agents of their respective activities), allows us to abstract out a new taxonomy, based on the behaviour rather than the general type of these entities



Fig. 1. A new 3-level abstraction hierarchy derived from verb/role combinations.

Figure 1 illustrates the creation of such a taxonomy, whose categories represent a yoking of verbs to specific case-roles, such as repair-agent and amend-agent, and whose category members are those HowNet concepts defined using these verbs and roles. The category-hopping nature of metaphor is now rather easily construed as a combination of generalization and re-specialization operations, in which one moves from one category to another by first passing through a common super-category like resume-agent. Thus, a surgeon can be seen as a repairman or a watchmaker, while a reviser of texts (an editor) can sometimes be seen as a surgeon. These metaphors make sense not because each is a human, but because each restores a better state.



Fig. 2. Newly derived HowNet categories may contain a diverse range of concepts.

Of course, this Aristotelian view of metaphor as an abstract “carrying-over” (the etymological origin of the word “metaphor”) can only be valid if concepts are ontologized by what they do, rather than by what they are (as is typically the case, in both WordNet and HowNet, and even Cyc [6]). Otherwise, metaphor could never

operate between semantically distant concepts, which it plainly does. For instance, figure 2 illustrates the derived taxonomy for HowNet concepts that are defined as agents of the verbs “kill”, “damage” and “attack”, each a specialization of the abstract verb MakeBad in HowNet. We see in this taxonomy the potential for famines to be metaphorically viewed as butchers and assassins, and for viruses to be seen as deadly intruders, or even man-eaters.

### 3 Deriving Fluid Category System

A taxonomy based on entity behaviour can harness the analogical potential of concepts in a way that can be exploited via simple generalization and re-specialization. However, such an approach begs a number of obvious questions about the nature of categorization that must first be answered. For instance, is every member of a category like kill-agent equally representative of that category? Is movement allowed between any two categories that share a common abstraction like MakeBad-agent, or is movement limited to certain members only, and in certain directions? When a concept moves from its conventional category to another, how is its degree of membership in this new category to be assessed? In this section we address the issue of fluid category structure, and in the following, we address the issue of directionality.

Every category will possess a prototype, a member that is highly representative of the category as a whole [5]. Such prototypes may be lexicalized; for instance, “killer” will be a highly representative of kill-agent, while the Chinese translation “杀手” is a composition of “killing” (杀) and “expert” (手). However, many categories like damage-agent have no obvious lexicalized prototype, so we need a more generic means of identifying the prototypical member of a category. Following Lakoff [5], the prototype will occupy a central position in the category’s structure, with other members organized in a radial fashion, at a distance from the centre that is inversely proportional to their similarity to the prototype. If we assume that the prototype will be that member that is most evocative of a category, we should first measure the evocation strength of each concept for a given category. This can be done by determining the frequency of occurrence of each concept within the category, and this, in turn, can be estimated by looking to a large corpus to see how each concept is actually employed by language users. Once the most evocative example is found for each category, membership scores can be assigned based on the strength of evocation. The corpus we use must be large, and while reasonably authoritative it must use words both literally and figuratively. For reasons outlined in section 5, we use here as our corpus the complete text of the open-source encyclopedia Wikipedia [11].

Thus, to estimate the membership level of the word-concept butcher|屠夫 in the category kill-agent, we first determine the corpus-frequency of the phrase “butcher who kills/killed”. In general, for estimating the membership of the concept C in the category V-agent, we use the query form “C who|which|that V”; for categories of the form V-instrument, we use the query “V with C”, and so on. Of course, some verbs are vaguer than others, and can have much higher corpus frequencies. We therefore need to normalize raw corpus-frequencies to obtain a truer picture of evocation

power. If  $f_{raw}(V\text{-role}: C)$  denotes the corpus frequency of concept  $C$  when considered as a member of the category  $V\text{-role}$ , where  $V$  is a verb like “kill” and  $role$  is one of agent, instrument, etc., then the adjusted frequency, a measure of true evocation, is estimated by:

$$f_{adj}(V\text{-role}:C) = \ln(f_{raw}(V\text{-role}:C)) \times \ln(\sum_x f_{raw}(V\text{-role}:x))^{-1} \quad (1)$$

Now, the prototype will be that member of a category with the strongest evocation:

$$Prototype(V\text{-role}) = \max_C(f_{adj}(V\text{-role}:C)) \quad (2)$$

The degree of membership of  $C$  in the category  $V\text{-role}$  is relative to the prototype:

$$Membership(V\text{-role}:C) = f_{adj}(V\text{-role}:C) \times f_{adj}(V\text{-role}: prototype(V\text{-role}))^{-1} \quad (3)$$

This ensures that the prototypical member has a membership score of 1, while all other members of a category will have a score in the range  $0 \dots 1$ . A concept can metaphorically be moved from a category in which it is conventionally a member to any other category in which it is considered to have a non-zero membership score, though as noted in [8], metaphor imposes some non-trivial constraints on directionality.

## 4 Directionality And Metaphor

These constraints on directionality of metaphor are not arbitrary, but go to the heart of metaphor’s cognitive function, which is to highlight aspects of a concept that are not considered salient within its conventional categorization. As Ortony [8] points out, one might describe a highway as a snake in order to emphasise the inherent danger of road travel. This sense of danger is highly salient in the category of snakes, but of low salience in the category of roads, so the move from latter category to the former has the effect of raising this salience. This salience imbalance means that most metaphors are asymmetric: the metaphor “ $X$  as a  $Y$ ” expresses a very different meaning to “ $Y$  is an  $X$ ”, and indeed, many metaphors cannot sensibly be reversed at all. Ortony’s theory of salience imbalance is computationally attractive because it provides a criterion for identifying metaphoric shifts even in propositions that are not semantically anomalous. For instance, one can meaningfully describe one’s lawyer as one’s bodyguard, to convey the idea that lawyers are protective of their clients. In contrast, if one describes one’s bodyguard as a lawyer, a very different meaning arises: since a lawyer is highly representative of the category argue-agent, while a bodyguard has but a weak membership in this category, we obtain instead a picture of a querulous bodyguard.

The fluidity of the category structures derived in section 3 means that a given concept may belong to a variety of different categories, with varying degrees of membership. Famine, for instance, belongs to the category damage-agent because that is where HowNet places it (insofar as HowNet defines famine via the verb “damage”). But it also belongs to the category kill-agent, since “famine that kills” is found 12 times in our corpus (and, as such is the most frequent verb for the subject famine);  $f_{adj}$ (kill-agent: famine) is 1.0. When one describes famine as a butcher or an assassin, one is implicitly placing famine in the category kill-agent, since butcher and assassin are highly representative of this category. In contrast, when one describes a virus as an intruder, the concept virus is re-categorized as an attack-agent, a category it sits well within, since  $f_{adj}$ (attack-agent: virus) is 0.4.

For the metaphor “X is Y” to be categorically well-formed, we hold that:

$$\text{cat}(X) = \max_{V\text{-role}}(f_{adj}(V\text{-role}:X)) \quad (4)$$

and

$$\text{cat}(Y) = \max_{V\text{-role}}(f_{adj}(V\text{-role}:Y)) \quad (5)$$

where

$$0 < f_{adj}(\text{cat}(Y):X) < f_{adj}(\text{cat}(Y):Y) \leq f_{adj}(\text{cat}(Y): \text{prototype}(\text{cat}(Y))) \quad (6)$$

In general, transparent, well-formed metaphors will maximize  $f_{adj}(\text{cat}(Y):X)$  and  $f_{adj}(\text{cat}(Y):Y)$  while obeying the above constraints, since  $f_{adj}(\text{cat}(Y):X)$  measures the degree of fit of the metaphoric topic in its destination category,  $\text{cat}(Y)$

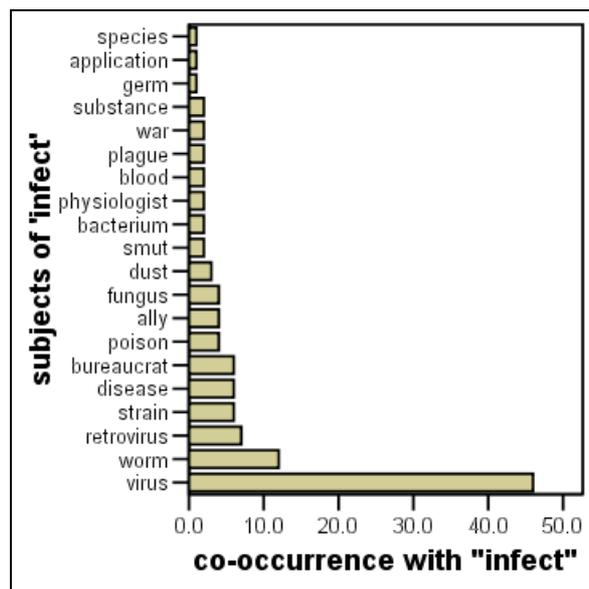
It is instructive to consider not only the fit of individual concepts within particular categories, but the general compatibility of categories amongst themselves. For any given pair of categories A and B, we can estimate the average membership of every member of A in B, and the average membership of every member of B in A. For the most part, these averages will not be the same, since the asymmetry of metaphor suggests that the likelihood of movement from A to B will be different than that for movement from B to A. The collected set of these inter-category movement potentials will form a slip-net of the kind described by Hofstadter [4], allowing a priori predictions to be made about category-shifting in general. This slip-net gives us some means of formalizing the surprise that derives from a strikingly original metaphor. For a particular metaphor that defies such predictions should therefore be seen as all the more surprising and creative, provided, of course that it also satisfies the constraints outlined above.

## 5 Preliminary Empirical Evaluation

The choice of corpus is clearly key to the quality of category-membership statistics that can be derived using the methods of sections 3 and 4. This corpus must be large,

it must be representative of language use in general, and it should offer a means of search that is robust in the face of noise. At first blush, then, the world-wide-web seems an ideal candidate: in size it is unmatched, and various APIs are available to access powerful search engines like Google. Unfortunately, such APIs rarely provide enough control over the query or the archive to ensure that noise can be eliminated, since these engines typically perform their own stemming and stop-word elimination, putting truly strict matching beyond our reach. This means that common noun-noun collocations, like “fossil record” and “share issue”, are easily confused for infrequent or nonsensical noun-verb collocations like “fossils that record” and “shares that issue”.

To ensure strict matching with controlled morphology, we require a local text corpus that we can index and search directly, and even subject to part-of-speech tagging. For this reason we choose the collected text of the open-source encyclopedia Wikipedia [11], which is available to download in XML form. Wikipedia has several obvious benefits as a text corpus: each document is explicitly tagged with a subject-label, since each article defines a specific headword; documents exist in a rich web of interconnections; and documents strive to be authoritative on their subjects. Consider the range of subjects that are found in Wikipedia for the verb “to infect” (with frequencies of co-occurrence shown in Figure 3):



**Fig. 3.** The range of subjects for the verb “infect”

Now consider the range of verbs that can be used with the subject “virus”:

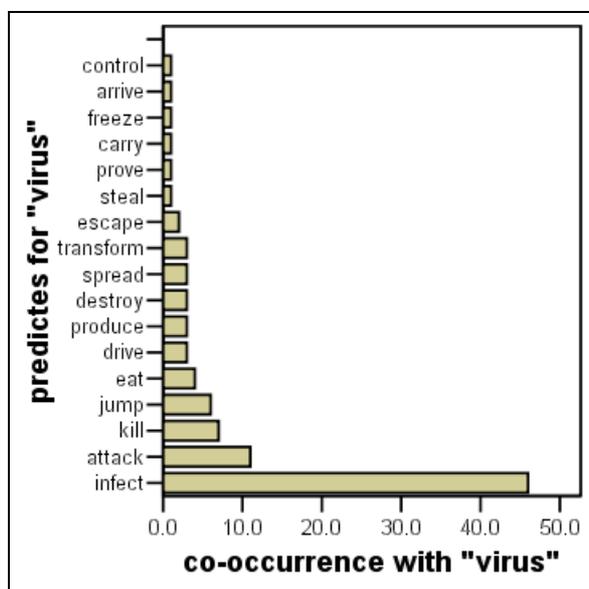


Fig. 4. The range of verbs for the subject “virus”

We see from these diagrams that Wikipedia contains enough diversity to capture the dominant application of each verb, and the dominant behaviour of each subject noun. Furthermore, Wikipedia contains enough diversity to reveal creative uses of these nouns and verbs; this snapshot reveals, for instance, that “smut” can “infect” (2 uses) and that a “virus” can “eat”, “escape” and even “steal”.

One can ask how well these corpus-derived category structures compare with the hand-crafted category structures of HowNet, since one can reasonably expect human-assigned category memberships to be a gold standard for this task. We find that in 69% of cases, the HowNet-assigned category for a given word-concept is also the dominant corpus-derived category, and that in 76% of cases, a word-concept has a statistical membership in the HowNet-assigned category that is greater than the median membership score for that category.

In fact, these results suggest that HowNet is far from being a gold-standard for category membership. In many cases, the HowNet category name is either poorly named or is dangerously misleading. For instance, the primary sense of the verb “doctor” in English is not “heal” but “fiddle” (as in “to doctor one’s résumé”). Likewise, HowNet assigns the name “resume” to the super-category of “repair” and “doctor”, when the verb “restore” is more appropriate in English. In many other cases, the HowNet assigned category is only one of several that seem intuitively appropriate. For instance, the word “knight” is assigned the dominant category protect-agent (based on 12 occurrences of the pattern “knight who protects”) while HowNet assigns it to the category defend-agent (which is the second-most popular corpus assignment, based on 10 occurrences of “knight who defends”). Viewed from this perspective, the corpus-based and hand-crafted approaches to category

assignment are complementary, not conflicting, where each can serve to validate and enrich the other.

## 6 Conclusion

Metaphor is a paradigmatic example of linguistic creativity at work. Its nature is both combinatorial – since metaphor operates through the juxtaposition of different concepts – and transformational, since these juxtapositions often alter our perception of category boundaries. In using corpus-frequencies to model and constrain the category-shifting character of metaphor, it may seem that we have inevitably reduced this creative potential to the level of Boden’s P-Creativity [1], insofar as only combinations that already occur in our corpus can be evaluated within this framework. Though true to an extent, this is not the complete picture. Since category membership is evaluated, and thus governed, by statistics collected from a large corpus, then the creativity of category membership is indeed limited to the level of P-Creativity only. However, for a metaphor of the form “X is a Y”, these statistics merely dictate the membership of concept X in the category  $\text{cat}(Y)$ . The choice of which particular concept Y to use, to evoke the category  $\text{cat}(Y)$ , is still very much a free choice on the part of the speaker, and is thus an opportunity for the speaker to contrive a H-Creative juxtaposition of concepts that has never been uttered before. Though much work remains to be done with this framework, we believe it is a promising foundation on which to study this most alluring of creative phenomena.

## References

1. Boden, M. (1999). Computational models of creativity. *Handbook of Creativity*, pp 351–373
2. Dong, Z. and Dong, Q. (2006). *HowNet and the Computation of Meaning*. World Scientific: Singapore
3. Glucksberg, S. and Keysar, B. (1993). *How Metaphors Work*. *Metaphor and Thought* (2nd edition), A. Ortony (Ed.). Cambridge University Press
4. Hofstadter, D. (1995). *Fluid Concepts and Creative Analogies*. Basic Books
5. Lakoff, G. (1987). *Women, Fire and Dangerous Things*. Chicago University Press.
6. Lenat, D. and Guha, R. V. (1990). *Building Large Knowledge-Based Systems*. Addison Wesley.
7. Miller, G. A. (1995). WordNet: A Lexical Database for English. *Communications of the ACM*, Vol. 38 No. 11.
8. Ortony, A. (1993). The role of similarity in similes and metaphors. *Metaphor and Thought* (2nd edition), A. Ortony (Ed.). Cambridge University Press.
9. Searle, J. (1993). *Metaphor*. *Metaphor and Thought* (2nd edition), A. Ortony (Ed.). Cambridge University Press.
10. Veale, T. (2005). Analogy Generation in HowNet. In proceedings of IJCAI’2005, the International Joint Conference on Artificial Intelligence.
11. Wikipedia open-source encyclopaedia: [www.wikipedia.org](http://www.wikipedia.org)