

Analogy Generation with HowNet*

Tony Veale

University College Dublin,
Creative Language Systems Group, Dept. of Computer Science,
Dublin 4, Ireland.
Tony.Veale@UCD.ie

Abstract

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

1 Introduction

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

With a well known lexical ontology like WordNet (see [Miller, 1995]), both structure-mapping and taxonomic approaches are problematic. The idea that a one-size-fits-all representation like WordNet will actually provide a hypernym like *kill-agent* seems convenient almost to the point of incredulity. As much as we want our ontologies to an-

ticipate future analogies with these pro-active categorizations, most off-the-shelf ontologies simply do not possess such convenient terms (see [Wong, 2004]). Similarly, WordNet lacks the propositional content that is the necessary grist for a structure-mapping approach. The semantic content that would ideally fill this role is not explicit, but implicitly resides in the unstructured textual glosses that annotate each lexical concept.

In this paper we explore whether another lexical ontology, the Chinese/English HowNet system (see [Dong, 1988; Carpuat *et al.* 2002; Wong, 2004]), is better suited to the demands of analogy generation. HowNet combines a taxonomic backbone with an explicit, if somewhat sparse and under-specified, propositional semantics. This combination allows us to evaluate the extent to which both structure-mapping and abstraction theories of analogy can be supported by the same lexical ontology.

2 Past Work

That analogy and metaphor operate across multiple levels of conceptual abstraction has been well known since classical times. Aristotle first provided a compelling taxonomic account of both in his *Poetics* (see [Hutton, 1982] for a translation), and computationalists have been fascinated by this perspective ever since. While the core idea has survived relatively unchanged, one must discriminate theories that apparently presume a static type-hierarchy to be sufficient for all abstraction purposes (e.g., [Fass, 1998]), from theories that posit the need for a dynamic type hierarchy (e.g., [Way, 1991; Veale, 2003]). One must also differentiate theories that have actually been implemented (e.g., [Fass, 1988; Veale, 2003,2004]) from those that are either notional or that seem to court computational intractability (e.g., [Hutton, 1982; Way, 1991]). Perhaps most meaningfully, one must differentiate theories and implementations that assume hand-crafted, purpose-built ontologies (e.g., [Fass, 1988]) from those that exploit an existing large-scale resource like WordNet (e.g., [Veale, 2003,2004]). The latter approach side-steps any possible charge of hand-crafting by working only with third-party resources, but at the cost of living with their perceived flaws and inadequacies.

Structure-Mapping theory is founded on the premise that the most satisfying analogies are those that operate at the

* The generous support of Enterprise Ireland is acknowledged.

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

By employing a single lexical resource, HowNet, to implement both the taxonomic abstraction *and* the structure-mapping theories of analogy, we have as a secondary goal a demonstration that both perspectives are not fundamentally opposed. Structure-mapping can be seen as a form of structural-abstraction, where one abstracts out the causal backbone of a concept, while taxonomic abstraction, if derived from the relational structure of a concept, can also be seen as a highly selective form of structure-mapping.

3 Comparing WordNet and HowNet

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

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

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

which can be glossed thus: “a surgeon is a human with an occupation in the medical domain who acts as the agent of a doctoring activity.” The {~} construct serves as a self-reference, to mark the location of the concept being defined in the given semantic structure. The oblique reference offered by the tilde serves to make the definition more generic, so that many different concepts can conceivably employ the same definition. Thus, HowNet uses the above definition not only for surgeon, but for medical workers in general, from orderlies to nurses, internists and neurologists.

Perhaps because HowNet relies less on hierarchical differentiation, it has a considerably less developed middle ontology than WordNet. For instance, most kinds of person in HowNet, from mathematicians to hobos, are placed di-

rectly under the hypernym *human*/人, eschewing the intermediate concepts like {*professional*}, {*specialist*} and {*worker*} that give substance to WordNet’s middle ontology. We note that HowNet does indeed define these concepts – but unlike WordNet, it does so at the leaf level where they add nothing to the internal structure of the taxonomy.

4 Analogy Via Relational Signatures

The skeletal nature of HowNet semantic definitions, combined with the wide-spread use of {~} as a generic reference, suggests how HowNet might support an efficient approach to analogical recall. By indexing each concept on a reduced form of its semantic definition – a *relational signature* – analogies will correspond to collisions between concepts with different definitions but with identical signatures. Such an approach can be efficiently implemented using simple string hashing of signatures, to detect analogical collisions between kitchens and factories, generals and admirals, ballet dancers and acrobats, or cruise missiles and arrows. The devil here is in the *lack* of detail: because HowNet’s definitions are frequently imprecise and fail to fully specify a concept, they allow others – potential analogues – to occupy the same reduced semantic space. The further we exacerbate this deficiency, indexing each definition on an increasingly diluted version of itself, the more distant and creative will be the analogies that are generated. For example, excluding the hypernym of a definition, or its domain markings, facilitates analogies between people and non-people, such as pests and persecutors, or hackers and viruses.

To implement both the abstraction and structure-mapping theories of analogy, we will explore the effectiveness of two kinds of relational signatures in the current work: atomic signatures based on taxonomic abstraction, and more template-like signatures based on generalized propositional content in which place-holder variables have been added.

4.1 Taxonomic Abstraction with HowNet

Given the general impoverishment of HowNet’s middle ontology (at least compared with that of WordNet), abstraction-based signatures should not be based directly on taxonomic organization. Rather, by instead deriving taxonomic signatures from the relational structure of a concept’s semantic definition, we can better capture the functional and behavioral nature of the concepts concerned. We can do this by focusing on that part of each definition that contains an explicit self-reference in the form {~}. For instance, consider the following semantic definition of *repairman*/修理工:

{*human*|人:HostOf={*Occupation*|职位},
{*repair*|修理:agent={~}}}

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

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

```

resume-agent
  repair-agent
    repairman|修理工
    watchmaker|钟表匠
  doctor-agent
    surgeon|医生
    herbalist|药农
  amend-agent
    reviser|修订者

```

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

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

```

AlterForm-instrument
  cut-instrument
    knife|刀
    razor|剃刀
  stab-instrument
    sword|宝剑
    lance|长矛
  split-instrument
    grater|擦菜板
    glasscutter|玻璃刀
  break-instrument
    scissors|剪
    chainsaw|油锯
  dig-instrument
    pickaxe|镐
    chisel|凿

```

Figure 2: a hierarchy of taxonomic signatures that facilitates analogy between instruments that “alter the form” of others.

Each concept is thus assigned two different taxonomic signatures: a direct signature ($F-role$) based on the specific relational structure of the concept, and another more general signature ($F'-role$) that is abstracted from this direct signa-

ture. These signatures effectively form an alternate taxonomy by which the lexical concepts in HowNet can be organized for analogical purposes. Figure 2 above illustrates a partial hierarchy derived from the HowNet semantics of various form-altering tools.

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

```

MakeBad-agent
  kill-agent
    assassin|刺客
    Death|死神
  attack-agent
    intruder|侵略者

```

Figure 3: diversity among concepts with the same signatures.

This diversity, as illustrated by Figure 3, means that the analogy “Death is an assassin” can be generated without recourse to a more abstract signature.

4.2 Structure-Mapping with HowNet

The structure-mapping approach also strives for abstraction, not through the selective creation of new taxonyms but through a form of structural *rarefaction*. Structure-mapping theory places particular emphasis on the causal backbone of a concept’s propositional content, which is usually projected unchanged from one domain to another ([Falkenhainer *et al.* 1989]). Based on this isomorphic alignment of relational structures, the entities contained in each structure are typically placed into a 1-to-1 correspondence with one another. The attributive modifiers of these entities play a more peripheral role in structure-mapping, but in approaches like *Sapper* [Veale and Keane, 1997] they often serve as a literal grounding for an analogy.

The semantic definitions provided by HowNet are already so skeletal and under-specified that we can operate on the assumption that they represent the relational backbone of a concept’s meaning. So in generating a set of structure-mapping signatures for a given concept, we will assume that each signature preserves the general form of a single proposition. Consider the HowNet definition of *blind person*|盲人:

```

{human|人:
  {disable|知道:
    OfPart={part|部件:PartPosition={eye|眼},
            whole={human|人}}
  experiencer={~},
  scope={look|看}}

```

In other words, a blind person has “a disability of the eye that affects one’s ability to look”. One finds precisely the same propositional structure in the HowNet definition of *lame*

person|拐子, except that *eye*|眼 is replaced with *leg*|腿 and *look*|看 is replaced with *walk*|走. The goal of a structure-mapping approach is to capture this semantic isomorphism while identifying *eye:leg* and *look:walk* as cross-domain counterparts. One way to do this is to generalize from each definition a signature that, by virtue of identity, signals a structural equivalence between individual definitions. For instance, the common structural signature for *blind person*|盲人 and *lame person*|拐子 might look like this:

$\{?:\{ill|病态:OfPart=\{?\},experiencer=\{\sim\},scope=\{?\}\}\}$

Generalized structural signatures like this can be generated using the following 7-step process:

1. Split each definition into multiple propositions, and generate a separate signature for each.
2. If a proposition describes a noun concept, replace its taxonomic head with a ? marker. In contrast, if a proposition describes a verb concept, replace its taxonomic head with its most specific hypernym.
3. Replace the conceptual arguments bound to each case-role of a predicate with the variable marker {?}. These markers will indicate positions in the signature where 1-to-1 correspondences between source and target structures can be made.
4. When a propositional sub-structure corresponds to the definition of another HowNet concept, replace the entire sub-structure with a {?} variable marker.
5. Replace predicates by their immediate hypernyms in the HowNet taxonomy. Thus, both *repair*|修理 in the definition of *repairman*|修理, and *doctor*|医治 in the definition of *surgeon*|医生, should be replaced by the hypernym *resume*|恢复 when generating their respective signatures.
6. Remove any explicit domain tag in a proposition from the corresponding signature (e.g., the assignment *domain*={*medical*|医} in the definition of *surgeon*|医生). This is necessary since analogy is meant to transcend domain boundaries.
7. Generalize any attributive value to its immediate hypernym. As such, step 3 above should not variablize the arguments of the attributive relations *modifier*, *manner*, *restrictive*, *host* or *content*.

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

$\{?:HostOf=\{?\}\}$
 $\{?:resume|恢复:agent=\{\sim\}\}$

[Note: because the *HostOf* relation always occurs with the binding *Occupation*|职位 in HowNet, only the latter signature is retained as an analogical index]

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

apostle|使徒
 $\equiv \{human|人:$
 $\quad \{believe|修理:$
 $\quad \quad agent=\{\sim\},$
 $\quad \quad content=\{humanized|拟人\},$
 $\quad \quad domain=\{religion|宗教\}\}$

person who knows inside story|个中人
 $\equiv \{human|人:$
 $\quad \{know|知道:$
 $\quad \quad agent=\{\sim\},$
 $\quad \quad content=\{fact|事情:$
 $\quad \quad \quad modifier=covert|隐秘\}\}$

These are also assigned the same structural signature:

signature $\equiv \{?:\{HaveKnowledge|有知:$
 $\quad \quad agent=\{\sim\},$
 $\quad \quad content=\{?\}\}$

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

5 Comparative Evaluation

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

We evaluate the taxonomic and structure-mapping approaches using four criteria: *coverage* – the percentage of unique HowNet definitions from which a valid signature can be derived; *recall* – the percentage of definitions for which at least one analogical counterpart can be found; *parsimony/precision* – the percentage of effective signatures that can actually be used to generate analogies (since a parsimonious approach will precisely generate just those signatures that are analogically useful); and *richness* – the complexity of the mappings involved, as measured by the average number of entity correspondences per analogy.

5.1 Evaluating Taxonomic Abstraction

5.1.1 Taxonomic Coverage

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

However, of the 23,507 unique definitions in HowNet, just 6430 employ this form of self-reference. The coverage offered by taxonomic signatures is therefore just 27% of the available definitions.

5.1.2 Taxonomic Recall

For these 6430 self-referential definitions, 1579 unique direct signatures are generated. In turn, another 838 abstract signatures are derived from these via predicate generalization. In total, 2219 unique taxonomic signatures are generated, revealing that in 8% of cases, the abstract signature of one definition corresponds to the direct signature of another.

A majority of these signatures (59%) serve to generate analogies for 6184 semantic definitions. The overall recall rate then is 26%. The most productive taxonomic signature is *control_agent*, which serves to analogically co-index 210 unique definitions.

5.1.3 Taxonomic Parsimony/Precision

Overall, 1,315 of all 2219 taxonomic signatures prove useful in co-indexing two or more definitions, while 904 taxonomic signatures are associated with just a single definition. The parsimony of the taxonomic approach is thus 59%.

5.1.4 Taxonomic Richness

Only one mapping, at the gross level of source and target concepts, can be generated by the taxonomic approach. For instance, the approach can recognize that *blind person*|盲人 and *lame person*|拐子 are analogous by virtue of sharing the taxonomic signature *disable-experiencer*. However, it cannot recursively determine the entity mappings *eye:leg* and *look:walk*. The taxonomic approach thus has a uniform mapping richness of 1.

5.2 Evaluating Structure-Mapping

5.2.1 Structure-Mapping Coverage

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

5.2.2 Structure-Mapping Recall

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

The most productive structural signature is:

{component.部分:whole={?}}

which serves to analogically co-index 397 unique semantic definitions.

5.2.3 Structure-Mapping Parsimony/Precision

With 79% of all structural signatures serving to index just a single definition, the parsimony of the structure-mapping approach must be judged as a low 21%.

5.2.4 Structure-Mapping Richness

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

5.3 Analysis of Results

The results of this comparison, as summarized in Table 1 below, force us to draw some important conclusions about the utility of HowNet for analogy.

	Taxonomic	Structure-Map	Combo
<i>Coverage</i>	.27	.90	.90
<i>Recall</i>	.26	.61	.72
<i>Parsimony</i>	.59	.21	.24
<i>Richness</i>	1	2.48	2.24

Table 1: Comparison of both approaches to analogy in HowNet

First, though the taxonomic approach is capped by the limited use of self-reference among HowNet definitions, it demonstrates a recall rate that closely approaches this ceiling, managing to find analogies of non-trivial complexity for 1 in 4 HowNet definitions. Because of its broader coverage, structure-mapping does considerably better, generating analogies for 3 in 5 definitions. A combination of both approaches (“combo” in Table 1) generates analogies for almost 3 in 4 definitions, which is most encouraging given the creative demands of analogy generation. This is especially so as we have considered here analogies between unique definitions, not unique words. A given definition-level analogy can be lexically realized in many, sometimes hundreds, of different ways.

6 Conclusions and Future Work

We conclude then that HowNet contains sufficient structure to realistically support both a taxonomic abstraction view and a structure-mapping view of analogy generation. Nonetheless, we need to investigate additional sources of semantic content to further increase the recall of the abstraction approach. For this, we must turn to the implicit content hidden in HowNet’s lexico-semantic structure.

Perhaps surprisingly, one source of implicit semantic content is orthography. Most Chinese entries in HowNet are multi-character – and thus multi-morpheme – terms whose composite orthography affords a kind of semantic transparency that other writing systems (e.g., that of English) do not possess. Thus, 手术刀, meaning “scalpel”, is a composite not just of characters but of ideas, for 手术 means “surgery” and 刀 means “knife”. Likewise, 哲学家, which translates as

“philosopher”, is a composition of 哲学 (“philosophy”) and 家 (“specialist” or “scientist”). In turn, *philosophy*|*哲学* is ontologized by HowNet as a specialization of *knowledge*|*知识*, as is *logic*|*辩学*, *mathematics*|*数学*, *lexicography*|*词典学* and even *midwifery*|*产科学*. By decomposing compound terms in this way, and by generalizing the extracted modifiers, a three-level taxonomy can be constructed to complement that which is formed by taxonomic signatures. From these examples alone, the partial taxonomy of Figure 4 can be derived:

knowledge-human
mathematics-human
 mathematician|数学家
philosophy-human
 philosopher|哲学家
midwifery-human
 midwife|产科
Buddhism-human
 Buddhist|佛教徒

Figure 4: Portion of an alternate three-level hierarchy derived from Chinese compound terms (prefix term + hypernym).

The analogical potential of such an alternative signature scheme becomes clear when one notices that it immediately supports Plato’s classical analogy of philosopher as midwife.

References

- [Carpuat *et al.*, 2002] Carpuat, M., Ngai, G. Fung, P., and Church, K.W. Creating a Bilingual Ontology: A Corpus-Based Approach for Aligning WordNet and HowNet. *The Proceedings of GWC 2002, the 1st Global WordNet conference*, Mysore, India.
- [Dong, 1988] Dong, Z. Knowledge Description: What, How and Who? *The Proceedings of the International Symposium on Electronic Dictionaries*, Tokyo, Japan.
- [Hutton, 1982] Hutton, J. *Aristotle's Poetics*. Norton.
- [Way, 1991] Way, E. C. Knowledge Representation and Metaphor. *Studies in Cognitive systems*. Kluwer Academic Publishers: Holland.
- [Miller, 1995] Miller, G. WordNet: A Lexical Database for English. *Comms. of the ACM*, 38(11).
- [Fass, 1988] Fass, D. An Account of Coherence, Semantic Relations, Metonymy, and Lexical Ambiguity Resolution. In: S. Small, G. W. Cottrell and M. K. Tanenhaus (eds.): *Lexical Ambiguity Resolution: Perspectives from Psycholinguistics, Neuropsychology and A.I.* Morgan Kauffman: CA.
- [Falkenhainer *et al.*, 1989] Falkenhainer, B.; Forbus, K.; and Gentner, D. Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence*, 41, pages 1-63.
- [Wong, 2004] Wong, S.H.S. Fighting Arbitrariness in WordNet-like Lexical Databases – A Natural Language Motivated Remedy. *The proceedings of GWC 2004, the 2nd Global WordNet conference*, Sojka, Pala, Smrz, Fellbaum, Vossen (eds.).
- [Veale and Keane, 1997] Veale, T. and Keane, M.T. The Competence of Sub-Optimal Structure Mapping on Hard Analogies. *The proceedings of IJCAI’97, the International Joint Conference on Artificial Intelligence*, Nagoya, Japan.
- [Veale, 2003] Veale, T. Dynamic Type Creation in Metaphor Interpretation and Analogical Reasoning: A Case-Study with WordNet. *The proceedings of ICCS2003, the 2003 International Conference on Conceptual Structures*, Dresden, Germany
- [Veale, 2004] Veale, T. WordNet sits the S.A.T.: A Knowledge-Based Approach to Lexical Analogy. *The proceedings of ECAI’2004, the 16th European Conf. on Artificial Intelligence*. John Wiley: London.