Qualia Extraction and Creative Metaphor in 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. There is an evident correlation between semantic distance and creativity in the treatment of metaphor. When the tenor and vehicle concepts of a metaphor are semantic neighbors, the local structure of the taxonomy can be used to constrain any interpretations and thus significantly curtail the breadth of the search space. This reliance on taxonomy makes WordNet ideally suited to the treatment of local metaphors. However, the distance involved in creative metaphors means that the tenor and vehicle rarely belong to the same semantic category, and interpretation must involve more than a simple recognition of a common taxonomic parent. Because creative metaphor effectively involves *reconceptualization*, interpretation must instead exploit the internal relational structure of the tenor and vehicle. In this paper we describe how certain key elements of the qualia structure of a concept, pertaining to its agentive and telic properties, can be automatically extracted from the sense glosses in WordNet.

1 Introduction

Conventional wisdom suggests that there is a strong correlation between the creativity of figurative comparisons like metaphor and the conceptual or semantic distance involved in the comparison. The suggestion goes something like this: the greater the conceptual distance between the concept being described (the *target* or *tenor*) and the concept carrying the description (the *source* or *vehicle*), then the greater the mental leap performed by the metaphor, and the greater the mental agility demanded of the listener. Creative metaphors are the high-wire circus acts of figurative speech, and are most impressive when performed without a safety net. Unlike the more mundane kind of metaphor, the kind that scarcely wanders outside of the category of the tenor to find a vehicle, the tension evoked by creative metaphor is caused primarily by the disparate nature of the concepts being combined, and the real danger that the combination may be misunderstood or rejected altogether as meaningless whimsy.

To appreciate figurative creativity, we should thus contrast what we might term 'local metaphor' with 'distant metaphor'. Because local metaphors involve tenor/vehicle combinations that belong to the same superordinate category (such as surgeons and butchers, cars and yachts or computers and rocketships), the structure of this superordinate can be used to tease out the small set of essential differences between both (see [1], [4], [5]). In contrast, distant metaphors involve an altogether more ambitious kind of category-hopping, and require an interpreter to recognize not the established taxonomic commonalities between both categories, which are too abstract to be useful, but the latent similarities between both that allow the tenor to be meaningfully reconceptualized in terms of the vehicle. Local metaphors do not challenge the established taxonomic ordering, whereas distant metaphors suggest new ways of seeing existing concepts, in terms of how they behave rather than how they are stored. As we shall see, both kinds of metaphor can be creative, but only the latter actually transcends the existing order as imposed by the central taxonomy.

Because we are interested in understanding the behaviors and affordances of lexical concepts, we adopt the generative view of lexical structure as championed by Pustejovsky [10]. We argue that in a generative lexicon we can expect to find precisely the kind of relational structure needed to interpret not only local metaphors like "a passport is a travel diary" (where both are travel documents), but also distant metaphors like "the protein is the magician of the cell". For practical reasons, our exploration of creative metaphor is conducted in the context of WordNet [3], a comprehensive lexical knowledge-base of English. The structure of WordNet makes explicit some of the relationships needed to construct a generative lexicon, most obviously the formal (taxonomic) and constitutive (meronymic) aspects of word meaning. But to truly test a computational model of metaphoric interpretation on a large-scale, it is necessary to augment these relationships with the telic and agentive components that are not encoded directly but merely alluded to in the textual glosses associated with each sense entry. In the sections to follow we describe a mechanism for automating the extraction of these relationships from WordNet glosses (in the same vein as [7]), and demonstrate how these extracted structures can be mapped and projected to generate appropriate interpretations for creative metaphors.

2 Qualia Extraction from WordNet Glosses

In a generative lexicon, the core elements of word meaning are represented by a nexus of relations called a *qualia structure*, which ties together the formal (i.e.,

hierarchical relations), constitutive (i.e., meronymic), telic (i.e., functional) and agentive (i.e., construction/creation) aspects of a word. For instance, a diary is *formally* a kind of 'book' that *constitutes* a 'collection of personal writings' whose *telic* purpose is to 'record' the observations of the *agent* that 'compiles' it. When a word like "diary" is used metaphorically, this relational nexus provides the structure for determining the internal systematicity of any interpretation. For instance, it is apt to describe a passport as a kind of travel diary since both are kinds of book (formal) that record (telic) travel experiences.

We describe here an approach to qualia extraction from WordNet glosses 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. Nevertheless, even this narrow slice yields a significant amount of qualia structure, since WordNet already encodes formal and constitutive relations in its taxonomic and meronymic links between synsets. We thus concentrate our efforts on the extraction of telic (i.e., goal-oriented) and agentive (activity-oriented) lexical relations.

We exploit the fact that the agentive and telic aspects of lexico-conceptual structure are often expressed using nominalized verbs that implicitly encode relational structure. A small number of highly productive morphology rules can thus be used to connect "observe" to "observer" and "observation" (and vice versa), "specialize", to "specializer" and "specialization", and so on. For example, the WordNet concepts *{botanist}* and *{philologist}* are both defined with glosses that explicitly employ the term "specializing", thus evoking the concept *{specializer}* (a hyponym of *{expert}*) Now, because *{specializer}* is compatible with the concepts *{botanist}* and *{philologist}* by virtue of being a hyponym of *{person}*, this in turn suggests that *{botanist}* and *{philologist}* should be seen as hyponyms of *{specializer}*, making *specializer_of* is an appropriate telic relation for each. Thus, using a combination of derivational morphology and simple taxonomic reasoning, the relational structure *specializer_of:specialization* can be associated with each concept. Since this structure is not already encoded in WordNet, it provides an additional dimension of similarity in any metaphoric mapping.

Broad clues as to the syntactic form of the gloss (such as the use of the passive voice) are also a valuable source of extraction information, especially when they can be robustly inferred from a simple combination of keyword analysis and inflectional morphology. For example, the passive voice should cause an extracted relation to be inverted, as in the case of *{dupe}*, whose WordNet gloss is *"a person who is swindled or tricked"*. The resulting relational structure is thus:

$\{dupe\} \rightarrow$ *of_swindler:swindler* $\wedge$ *of_trickster:trickster*

This approach requires the WordNet taxonomy to act as a vital sanity-check for any extracted relationship. In general, it is sensible to associate a relation *r* with a concept *c* if the nominalization of *r* denotes a concept that belongs to the same taxonomic category as *c*; thus, it is sensible to ascribe a *specializer_of* relation to *{botanist}* only because *{specializer}* and *{botanist}* each specify a sub-category of *{person}*. However, this broad injunction finds an important exception in metonymic contexts. Consider the WordNet gloss for *{diary, journal}*, *"a daily record of (usually private) experiences and observations"*, which yields the extracted relationships *of_diarist:diarist*, *of_experience:experience*, *recorder_of:recording* and *observer_of:observation*. A taxonomic sanity-check reveals that *{diary, journal}*, as a sub-category of *{communication}*, is not compatible with either *{recorder}* or *{observer}*, both sub-categories of *{person}*. However, it is taxonomically compatible with the objects of these relations, *{recording}* and *{observation}*, which suggests that a diary is both the object of, and a metonym for, the diarist as observer and recorder. This metonymy is most evident in the familiar address "dear diary", in which the diary is conceived as a personified counterpart of the observer. The concept *{diary, journal}* therefore yields the modified relational structure:

$\{diary, journal\} \rightarrow$ **observer_of:observation* $\wedge$ **recorder_of:recording* $\wedge$ *of_experience:experience*

The (*) here signals that the *observer_of* and *recorder_of* relations hold metonymically rather than literally. The presence of these relationships facilitate creative uses of the concept *{diary}* that follow the general pattern whereby artifacts are viewed from an intentional stance. For example, consider that the WordNet gloss for the concept *{witness, spectator}* is "a close observer", so that the following relational structure is extracted:

$\{witness, spectator\} \rightarrow$ *observer_of:observation*

It now becomes apt to metaphorically consider a diary to be a witness to one's life experiences. In structure-mapping terms, this aptness is reflected in the internal systematicity of finding a key relationship, *observer_of:observation*, common to each of *{diary}* and *{witness, spectator}*.

3 Local Metaphor

Local metaphors are juxtapositions of a source and target concept that can be united via a common taxonomic parent. Such metaphors are local since, by necessity, this common parent cannot be very far removed from the source and target or the interpretation will be too vague to be meaningful. Thus, if the common parent is conceptually close to the source and target, then the source and target must already be conceptual neighbors in the taxonomy.

Some advocates of this taxonomic approach allow the common parent to carry the meaning of the metaphor as a whole (e.g., see [4]). Thus, butcher is a meaningful metaphor for surgeon because both are professionals that cut flesh for a living. However, if we place the burden of meaning on the common parent, then by definition such metaphors will be symmetric in nature. Symmetric metaphors do exist, but they have little semantic tension and do not rise above the status of simple simile; for example, consider "credit unions are (like) banks" and "gamblers are [like] alcoholics". In contrast, the most creative metaphors are asymmetric [8], since they impose the highly-developed relational structure of the source concept onto that of the less-developed target [6, 9]. Without this imposition of relational structure, metaphor can be used only to highlight existing similarities rather than to actually create new ones, and is thus robbed of its creative function.

Notwithstanding this perspective, local metaphors do not have to be either symmetric or lacking in creativity. The key to seeing local metaphor as creative is to move the burden of meaning from the common taxonomic parent onto the structure that is projected from the source to the target. This 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 retrieve a counterpart for the source using the target as a retrieval cue, and then project the relational structure of the source onto this counterpart [6]. For instance, WordNet contains a variety of concepts that are formally similar to *{diary, journal}* and which also mention "travel" in their glosses, such as *{travel_guidebook}* and *{passport}*. These facilitate the following reading:

"travel" + {*diary, journal*} →
{*passport*} + **observer_of:travel:observation*
∧ **recorder_of:travel:recording*
∧ *of_experience:travel:experience*

Projecting the relational structure of *{diary, journal}* onto *{passport}* causes the latter to be seen as a journal of travel observations and experiences. Indeed, many travelers retain old passports for this very purpose.

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-category of *{person}*. But to fully appreciate why *{surgeon}* is more apt than other hyponyms of *{person}*, like *{astrologer}* say, one must look to the shared relational structure that is highlighted by the metaphor. 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}* thus become associated with the relationship *specializer_of:specialism*. This common relational structure facilitates the measurement of what we have termed 'internal systematicity' (in the Gentner sense). Thus, *{surgeon}* is seen as an apt vehicle for *{pastry_cook}* as both are people that specialize in a particular field. Instantiation of the shared structure leads to the following interpretation:

"pastry" + {*surgeon*} → {*pastry_cook*} + *specializer_of:* *pastry*:*surgery*

The choice to delve deeper, and recursively determine an appropriate interpretation of "pastry surgery", is left to the comprehender, who may instead choose to read the metaphor as a simple request to view pastry chefs as specialists. But this raises the question of how much structure must be shared for an interpretation to appear apt rather than merely inept. For example, one can equally well say "pastry linguist" or "pastry geologist" to highlight the specialist nature of pastry chefs, since *{geologist}* and *{linguist}* are also associated with an extracted *specializer_of* relationship. What makes these alternate metaphors seem clumsy is the difficulty in assigning appropriate interpretations to the recursive metaphors that they imply: "pastry geologist" implies the metaphor "pastry geology", while "pastry linguist" implies the metaphor "pastry linguistics".

(?) *"pastry"* + {*linguist*} → {*pastry_cook*} + *specializer_of:*pastry*:linguistics*

Whereas "pastry linguistics" is a non-starter in WordNet terms, "pastry surgery" has more potential for meaningful interpretation:

{*plastic_surgery*} ≡ *surgery concerned with therapeutic or cosmetic reformation of tissue*

"pastry" + {*surgery*} → {*plastic_surgery*}
+ *reformation_of: pastry:tissue*

This interpretation requires that an existing form of surgery is recruited and adapted so as to accommodate the concept *{pastry}*. In taxonomic terms, *{plastic_surgery}* is perhaps most appropriately adapted for this purpose, since *{tissue}* and *{pastry}* are both hyponyms of *{substance}* in WordNet.

It is often possible to recruit an established metaphor schema (see [9]) to accommodate an interpretation, if this schema has already been lexicalized in WordNet. Consider for instance the metaphor "genetic cartographer":

{*cartographer*} ≡ *a person who makes maps*
{*geneticist*} ≡ *a person who specializes in genetics*

∴ *"genetic"* + {*cartographer*} → {*geneticist*} + *mapper_of: genetic:mapping*

Here is it possible to interpret the sub-metaphor "genetic mapping" in terms of an existing metaphor in WordNet:

{*chromosome_mapping*} ≡ *the process of locating genes a chromosome*

A "genetic cartographer" is thus a geneticist that performs chromosome mapping.

4 Distant Metaphor

Distant metaphors involve a juxtaposition of source and target concepts that are considerably removed from each other in the taxonomic scheme of things. If such pairings have a common taxonomic superordinate, it is generally too abstract or too vague to contribute much in the way of structure or constraints to the meaning of the metaphor. Distant metaphors thus lack the safety net of a well-developed common ground against which their meaning can be determined.

For example, consider the metaphor "the protein is the magician of the cell". This metaphor juxtaposes two concepts, *{protein}* and *{magician, conjuror}*, whose lowest common hypernym in WordNet is *{entity, physical_thing}*. Since this hypernym is a unique beginner with tens of thousands of hyponyms, very little is gained by using this concept as a common ground for the metaphor. Rather, the metaphor must be understood in terms of the structural relations imposed on *{protein}* by the vehicle *{magician, conjuror}*:

$$\{magician, conjuror\} \equiv someone\ who\ performs\ magic\ tricks\ to\ amuse\ an\ audience$$
$$\rightarrow performer\_of{:}magic\_trick \wedge entertainer\_of{:}entertainment$$

$$\therefore \text{“protein”} + \{magician, conjuror\} \rightarrow performer\_of{:}\ \underline{protein}{:}magic\_trick$$
$$\wedge\ entertainer\_of{:}\ \underline{protein}{:}entertainment$$

By projecting the relational structure of *{magician, conjuror}* (as extracted from its WordNet gloss) onto *{protein}*, we succeed in shifting the onus of interpretation onto two sub-metaphors: “protein magic trick” and “protein entertainment”, whatever these might mean. However, each of these sub-metaphors might in turn yield to analysis as either a local or a distant metaphor. We consider here a local analysis in taxonomic terms, and rephrase the problem thus: since *{magic_trick}* and *{entertainment}* are types of activity in WordNet, what other hyponyms of activity can be found that involve proteins? To find any such hyponyms, we need to construe activity in its widest possible sense, rather than confine our search to the specific sense entailed by *{entertainment}* and *{magic_trick}*. This broadening of senses, a kind of ‘domain incongruence’ (see [11]), is often necessary in distant metaphors since the relational structure of one concept is being used to restructure a concept from a radically different taxonomic tree. So construing “activity” in its sense of *{activity, natural_process}*, we find two candidate hyponyms:

$$\{active\_transport\} \equiv transport\ of\ a\ substance\ (as\ a\ protein\ or\ drug)\ across\ a\ cell\ membrane\ against\ the\ concentration\ gradient$$

$$\{proteolysis\} \equiv the\ hydrolysis\ of\ proteins\ into\ peptides\ and\ amino\ acids\ by\ cleavage\ of\ their\ peptide\ bonds$$

At this point we reach the limits of a WordNet-powered analysis, as it is a matter of subjective opinion as to whether either of these two activities can be considered magical or entertaining. We are limited instead to rather superficial measurements of aptness, according to which, *{active_transport}* seems the most apropos candidate, if only because it explicitly mentions the thematic keyword “cell”. We can thus paraphrase the WordNet interpretation of “the protein is the magician of the cell” as follows: in facilitating active transport in a cell, a protein acts in a way that can be considered both magical and entertaining.

5 Conclusions

We conclude with some empirical observations regarding the efficacy of extracting qualia structures from WordNet glosses. The process as it is described here, operating entirely via a combination of derivational morphology rules and taxonomic sanity-checking, is currently able to automatically extract relationships from 40% of the noun glosses in WordNet 1.6. Furthermore, 96% of all noun glosses contain at least one word with a denotation in this extraction set, which suggests that future extensions to the process may be able to obtain much higher coverage with relatively minor additions to the mechanism. We are encouraged in this view by recent events in the development of WordNet, which will soon be augmented with explicit, hand-coded morpho-semantic connections between verbs and their nominalizations (see [10] for news of these developments). This hand-coding should increase the precision of the extraction process while entirely removing its dependency on over-generating morphology rules.

WordNet is not an encyclopedia (in the sense of [2], say), but a lexical taxonomy with some of the useful features of an encyclopedia. This essential difference has both advantages and disadvantages. On the plus side, the brevity of the textual definition associated with each WordNet synset means that this gloss contains highly relevant information; were our extraction techniques to be applied to much larger definitions, the quality of the extracted relationships would surely be lowered dramatically. On the negative side, this means we can expect to extract just a sliver of each concept's useful relational content. To properly handle the recursive sub-metaphors that are entailed by creative metaphors (such as viewing the activity of a protein as an entertaining magic trick), we will need to draw upon a richer source of word and world knowledge. To this end, it may be worthwhile to consider the usefulness of large text corpora, and perhaps even the text content dynamically available on the World Wide Web, as a basis for interpreting creative metaphors.

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)
11. Tourangeau, R., Sternberg, R. J.: Aptness in Metaphor, Cognitive Psychology 13 (1981).