Systematicity and the Lexicon in Creative Metaphor

Tony Veale

Department of Computer Science,
University College Dublin, Belfield, Dublin 6, Ireland.
Tony.veale@UCD.ie

Abstract

Aptness is an umbrella term that covers a multitude of issues in the interpretation and generation of creative metaphor. In this paper we concentrate on one of these issues — the notion of lexical systematicity — and explore its role in ascertaining the coherence of creative metaphor relative to the structure of the target concept being described. We argue that all else being equal, the most apt metaphors are those that resonate most with the way the target concept is literally and metaphorically organized. As such, the lexicon plays a key role in enforcing and recognizing aptness, insofar as this existing organization will already have been lexicalized. We perform our exploration in the context of WordNet, and describe how relational structures can be automatically extracted from this lexical taxonomy to facilitate the interpretation of creative metaphors.

1 Introduction

When one considers the aptness of creative metaphor and how one might measure it, one finds a whole range of issues lurking between the apparent unity of this umbrella term. This

complexity is compounded by the fact that metaphors operate at several different levels of representation simultaneously: the conceptual level, or the level of ideas; the lexical level, or the level of words; and the pragmatic level, or the level of intentions. A metaphor may fall at any of these hurdles, either through a poor choice of a source concept, a poor choice of words in communicating this concept, or in a failure to observe the expectations of the context in which the metaphor is expressed.

Some degree of aptness is afforded by metaphors that compare semantic neighbors, inasmuch as the existence of a common taxonomic parent suggests that the source and target are in the same, or at least similar, domains (e.g., see Way, 1991). For instance, metaphors that compare politicians to architects, or even geneticists to cartographers, derive some measure of aptness from the fact that in each case the source and target are sub-categories of the *Profession* category. However, since the most creative of metaphors are those that make the greatest semantic leaps between the source and target concepts, such category-hopping metaphors do not have the luxury of comparing concepts that are already deemed similar in taxonomic terms, as evidenced by a common superordinate concept, but must instead establish a new basis for conveying similarity that is not itself taxonomic. Consider for

instance a corollary of the above metaphor in which “genomes are maps”. The aptness of these similarity-creating metaphors is instead a measure of the isomorphism between the relational structures of the source and target, so that the concepts with the greatest structural overlap will often produce the most apt metaphors. In this respect, metaphoric aptness is a function of what Gentner terms the *systematicity* of a structure-mapping. According to (Gentner, 1983) and the structure-mapping school of thought (e.g., see also Veale and Keane, 1997), the best interpretations of a metaphor or analogy are those that systematically pair-off the greatest amount of connected relational structure in each concept. We refer to this kind of structural aptness as *internal systematicity*, since any sense of aptness arises out of a coherence between the internal structures of the concepts being mapped.

Lakoff and Johnson (1980) also place a strong emphasis on metaphoric systematicity, but in their hands the notion is construed in more *external* terms. To L&J, systematicity is a measure of the generativity of a metaphoric schema, so that the same schema (such as *Life is a Journey*) can serve as the deep structure for a wide variety of different, but mutually systematic, surface metaphors (such as “my job has hit a rocky patch” and “my career has stalled”). In this view, systematicity is a measure of how much a metaphor resonates and coheres with existing metaphors for thinking about the target concept, so that when viewed collectively, they together suggest the operation of a common underlying schema. This view of systematicity is external to the concepts involved since it predicates their aptness to each other on the existence of other structures (metaphor schemas) into which they can be coherently connected.

In this paper we argue that the lexicon is central to the determination of both kinds of systematicity, internal and external, especially if one is an adherent of the *generative lexicon* view of word meaning as championed by (Pustejovsky, 1991). In such a lexicon we can expect to find precisely the kind of relational structure needed to perform structure mapping and thereby measure the internal systematicity of a metaphor like “a passport is a travel diary”. In addition, we can expect to find the

lexicalized metaphor structures that represent the surface manifestations of existing modes of thought, and it is against these structures that the external systematicity of an interpretation can be measured.

This research is conducted in the context of WordNet (Miller, 1995; Fellbaum, 1998), 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 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 (in the same vein as (Harabagiu *et al.* 1999), and for using them to generative apt interpretations for metaphors involving WordNet entries.

2 Qualia Extraction from 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 rules¹ 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:

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

Note that the extraction process is too shallow to do very much with the disjunctive “or” present in the gloss of *{dupe}*, as this is more a process of information extraction than full natural-language parsing. Thus, it simply conjoins any relationship that can be reliably extracted with morphological cues into an overall relational structure. This structure is simply a bag of relations at present, which we choose to present here as connected via conjunction. Future versions of the extraction process may attempt to impose a more elaborate inter-connecting structure on the relationships that are extracted, but for the present, an unstructured bag is sufficient to support a consideration of metaphor in WordNet.

Since morphology alone is not a sufficiently reliable guide for extraction purposes, the approach crucially 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:

¹ The developers of WordNet have recently announced that hand-coded morpho-semantic links will be added to future versions of WordNet, to make explicit the relationship between verbs and their nominal forms, thus obviating the need for such rules while making the extraction task even more reliable.

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

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 the concepts *{diary}* and *{witness, spectator}*.

3 Internal Systematicity

Because purely taxonomic interpretations are created on the basis of commonalities, they tend to be highly symmetric, as in the case of similes such as “credit unions are like banks” and “gamblers are like alcoholics”. In contrast, the most creative metaphors are asymmetric (Ortony, 1991), since they impose the highly-developed relational structure of the source concept onto that of the less-developed target (see Lakoff and Johnson, 1980; Gentner, 1983; Veale and Keane, 1997). 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.

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 (Gentner, 1983). 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}*.

$$\begin{aligned} “travel” + \{diary, journal\} &\rightarrow \\ \{passport\} + *observer_of:travel:observation & \\ \wedge *recorder_of:travel:recording & \\ \wedge of_experience:travel:experience & \end{aligned}$$

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.

Metaphors are most apt when projection highlights a latent relational structure that already exists in the target concept (Ortony, 1979). 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:

$$\begin{aligned} “pastry” + \{surgeon\} &\rightarrow \\ \{pastry_cook\} + specializer_of:pastry:surgery & \end{aligned}$$

One can reasonably argue that much more sophisticated interpretations are available to human readers of this metaphor, e.g., that pastry cooking and surgery are both delicate operations

involving special training, both are performed with specialized instruments in very clean surroundings, etc. But given the inherent limitations of working with an existing semi-structured knowledge source such as WordNet, as opposed to a dedicated, hand-crafted knowledge-base, “pastry specialist” must suffice as a generalization for these richer interpretations. Alternately, one might argue that it is “pastry” rather than “surgeon” that undergoes metaphoric reinterpretation, so that the phrase denotes a literal surgeon that operates on metaphoric pastries, such as movie starlets or supermodels. In this current work we choose to focus on the relational potential for the head word to metaphorically denote a relationally similar, if sometimes semantically distant, referent, while acknowledging that this illuminates just one part of the picture.

Nonetheless, interpretations like “pastry specialist” can be given more credibility if one delves deeper into its metaphoric ramifications to consider the recursive sub-metaphors that it implies. For instance, as stated in the analysis above, “pastry surgeon” implies the plausibility of a meaningful interpretation for “pastry surgery”. This 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

There is little that can be done to put a sensible interpretation on “pastry linguistics” in WordNet, given the taxonomic and relational structure of *{pastry}* and *{linguistics}*. In contrast, “pastry surgery” has more potential for meaningful interpretation using WordNet structures. There exists a sense of surgery that denotes a discipline in the natural sciences, and from *{pastry}* a broad search will find the concept *{dietetics}*, another discipline of the natural sciences dedicated to food preparation. This analogue of *{surgery}* can be found by first considering all concepts associated with “pastry”, then all concepts associated with “baked goods”, then “foodstuff” and “food”, until an appropriately similar candidate is found.

{dietetics} ≡ *the scientific study of food
preparation and intake*

This is not a particularly well-known concept, so it would be difficult to argue that this forms the cornerstone of an easily understood metaphor like “pastry surgeon”. However, the concept *{dietetics}* does at least concretize, in WordNet terms, the idea that one can take a precise, scientific view of food preparation, and it is the plausibility of this notion that allows us to make sense of pastry preparation as a surgical activity. There is no true substitute for situated experience of the world, but when it comes to metaphor interpretation using lexical resources like WordNet, we should be willing to use any lexical precedent we can find.

As an alternate strategy, we can seek to recruit a sub-category of surgery that can be modified in some way to accommodate the concept *{pastry}*. One such category is *{plastic_surgery}*, whose gloss reveals a concern with the reformation of body tissue.

{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. Of course, the intended sense of “tissue” in the above gloss is not *{tissue}*, *tissue_paper}* but *{tissue}* as a hyponym of *{body_part}*. However, creative metaphors often involve a degree of domain incongruence, whereby a given word has a different meaning in the source and target domains (Ortony, 1979). In fact, one might say that domain incongruence is essential to creative metaphor, since interpretation will necessitate the grafting of structure from radically distant parts of the concept ontology, and such grafts may fail if the features involved maintain their strict, source-dependent definitions.

4 External Systematicity

Metaphors appear more apt when they systematically evoke, or connect into, established modes of metaphoric thought. This is systematicity considered from an external vantage as described by (Lakoff and Johnson, 1980). For example, when processing the metaphor “political mechanic”, several concepts can be reached from “political” that prove to be taxonomically compatible with *{mechanic}*, among them *{political_leader}*, *{political_scientist}* and *{machine_politician}*. However, closer inspection of the projected structure suggests that the last, *{machine_politician}*, is the most systematic:

“political” + {mechanic}
 $\rightarrow$ *{machine_politician}*
 $+ \text{machinist_of: } \underline{\text{political:machine}}$

Because the extracted qualia structure for *{mechanic}* hinges on the relationship *machinist_of:machine* there is a suggestive lexical systematicity with the concept *{machine_politician}*. Furthermore, the instantiated structure creates a fortuitous pairing *political:machine*, which already exists in WordNet as the lexicalized metaphor *{political_machine}*. This marks “political mechanic” as a systematic outgrowth of the established metaphor schema *Political System As Machine* (whose corollary is *Political Operatives*

as *Fixers*). The same schema comes into play when interpreting the metaphor “political draftsman”, whose WordNet gloss also evokes images of machinery.

Lexicalized metaphors like *{political_machine}*, *{political_scientist}* and *{political_campaign}* act as the recognizable landmarks in the search space of possible interpretations for novel metaphors. So if an interpretation can be generated that connects into an established metaphor, it has a greater provenance than one that stands alone. Here are some further examples:

{torchbearer} $\equiv$ *a leader in a campaign or movement*
 $\therefore$ *“political” + {torchbearer}* $\rightarrow$
{political_leader} + campaigner_of: political:campaign

{missionary} $\equiv$ *someone who attempts to convert others to a [...] program*
 $\therefore$ *“political” + {missionary}* $\rightarrow$
{political_commissar} + programmer_of: political:program

{sociologist} $\equiv$ *a social scientist who studies [...] human society*
 $\therefore$ *“political” + {sociologist}* $\rightarrow$
{political_scientist} + scientist_of: political:science

These examples are fortuitous in the sense that the instantiation of qualia structure directly suggests an existing WordNet concept. In most cases, however, the external systematicity becomes visible only upon recursive consideration of the instantiated structure as a source of metaphor in itself. Consider the metaphor “genetic cartographer”, for which *{geneticist}* is retrieved as a thematically similar concept:

{cartographer} $\equiv$ *a person who makes maps*
{geneticist} $\equiv$ *a person who specializes in genetics*
 $\therefore$ *“genetic” + {cartographer}* $\rightarrow$
{geneticist} + mapper_of: genetic:mapping

There is no denotation for “genetic mapping” in WordNet, so at first blush the above interpretation

fails to connect into an existing lexicalized metaphor. However, when we recursively consider the combination “genetic mapping” as a metaphor in itself, we obtain the following interpretation:

$$\begin{aligned} \text{“genetic”} + \{\text{mapping}\} &\rightarrow \{\text{chromosome_mapping}\} \\ &\equiv \text{the process of locating genes} \\ &\quad \text{on a chromosome} \end{aligned}$$

This allows us to recognize “genetic mapping” as an alternate way of denoting the concept $\{\text{chromosome_mapping}\}$, while the fact that a mapping metaphor has already been lexicalized in the genetics domain allows us to recognize the external systematicity inherent in the interpretation of “geneticist as cartographer”. This WordNet entry serves to ground the sub-metaphor of genetic mapping in an existing concept, allowing the recursive analysis of sub-metaphors to halt at this point. A “genetic cartographer” is thus a geneticist that performs a specialized kind of map-making called *chromosome mapping*, where the terrain that is mapped is biological and information-theoretic rather than geological or geographic. Though chromosome mapping is itself a metaphor, its independent existence in WordNet means that it does not need to be justified in the context of an interpretation of “genetic cartographer”, and for the purposes of analysis can be treated as a literal stopping-point.

5 The Challenge of Aptness

I suspect we can all agree that aptness involves a complex interaction of different issues that arise from lexical and conceptual choice. The real question is the degree to which each of these issues influences a particular interpretation, and the weighting, if any, that is to be given to each component of aptness in an algorithmic model. Take the metaphor “political surgeon”: by considering the concepts in the semantic neighborhood of $\{\text{surgeon}\}$ reachable via the thematic cue “political”, we find the following competing interpretations:

$$\{\text{political_scientist}\} \equiv \text{a social scientist } \underline{\text{specializing}} \text{ in the study of government}$$

$$\{\text{spin_doctor}\} \equiv \text{a spokesperson for a political party or candidate who tries to forestall negative publicity}$$

The first of these interpretations, $\{\text{political_scientist}\}$, is apt for reasons of internal systematicity, as both it and $\{\text{surgeon}\}$ have an extracted qualia structure that contains a *specializer_of:specialization* relationship. This leads to the following interpretation:

$$\begin{aligned} \text{“political”} + \{\text{surgeon}\} &\rightarrow \{\text{political_scientist}\} \\ &+ \text{specializer_of:}\underline{\text{political}}\text{:specialization} \end{aligned}$$

The second interpretation, $\{\text{spin_doctor}\}$, does not exhibit the same internal systematicity, but it does exhibit an external systematicity of sorts: the head of this compound term, “doctor”, denotes a concept $\{\text{doctor, physician}\}$ that is a hypernym of the metaphoric vehicle, $\{\text{surgeon}\}$.

It would seem a matter of personal choice as to which interpretation should be privileged here, as different listeners may attach more weight to the presence of internal systematicity in $\{\text{political_scientist}\}$ than to the suggestion of external systematicity in $\{\text{spin_doctor}\}$, and vice versa. This suggests that the problem of aptness determination involves a great deal of hidden parameters yet to be made explicit in any model of interpretation. As researchers interested in computational treatments of metaphor, our goal then should be to explicate what factors we can in algorithmic and representational terms, to provide the basic inventory of components needed to proceed with our investigation into this elusive and considerably vexing phenomenon. In this paper we have argued that the natural place to compile this inventory is the lexicon, since this acts as the bridge between word and world knowledge and aptness is a phenomenon that hops freely between both.

References

- George Lakoff and Mark Johnson. 1980. *Metaphors we live by*. University of Chicago Press: Chicago.
- George A. Miller. 1995. WordNet: A Lexical Database for English. *Communications of the ACM*, Vol. 38 No. 11.
- Christiane Fellbaum. 1998. *WordNet: An electronic lexical database*. Cambridge, England: Cambridge University Press, Cambridge, MA.
- Eileen C. Way. 1991. Knowledge Representation and Metaphor. *Studies in Cognitive systems*. Kluwer Academic Publishers, Amsterdam.
- James Pustejovsky. 1991. The generative lexicon. *Computational Linguistics*, Vol. 17 No. 4.
- Dedre Gentner. 1983. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, pp 155-170.
- Tony Veale and Mark T. Keane. 1997. The Competence of Sub-Optimal Structure Mapping on Hard Analogies. *The proceedings of IJCAI'97, the International Joint Conference on Artificial Intelligence, Nagoya, Japan*. Morgan Kaufman, San Mateo CA.
- Sanda Harabagiu, George A. Miller, and Dan Moldovan. 1999. WordNet 2 - A Morphologically and Semantically Enhanced Resource. *The Proceedings of the ACL SIGLEX Workshop: Standardizing Lexical Resources*. Maryland, USA.
- Andrew Ortony. 1979. The role of similarity in similes and metaphors. In Ortony, A. (ed.): *Metaphor and Thought*. Cambridge University Press : Cambridge, U.K.