Metaphors Break the Rules

Michael Witbrock, Douglas Lenat

09/09/09

Corresponding Author:

Dr. Michael Witbrock
witbrock@cyc.com
Cycorp Inc, Suite 200
7718 Wood Hollow Drive, Austin, Texas, 78731
Phone: 512 342 4003
Fax: 512 342 4040

1. Summary

Metaphor is ubiquitous in natural language, and is highly functional, not just decorative. Unfortunately, even the latest encouraging research results in automating its analysis fall well short of the *understanding* of metaphor required to apply it to cultural analysis. Their coverage is too limited, and they require expensive training data.

However, it is a general feature of metaphors that they violate hard or soft semantic constraints. They can be analyzed by detecting that these constraints have been violated, and describing how. Some of the gap in understanding can be closed simply by having a crisp model of argument types for the predicates that are denoted in text – a violation of those types is evidence for a metaphor. Closing the remaining gap, for more ambiguous cases, will require discourse analysis and real world knowledge to eliminate interpretations that are conceivable, but very unlikely. This sort of semantic analysis of metaphor is both possible and practical, and can be done at relatively modest cost by exploiting the state of the art in knowledge bases, inference and scalable shallow text analysis.

Beyond detection, metaphors can be classified semantically. Recognizing what kind of thing has violated a constraint, and the nature of the constraint itself, serves to identify a metaphorical transformation. For example "Wal-Mart rooted itself firmly in the US before spreading its branches overseas" is an example of the "ORGANIZATIONS ARE PLANTS" metaphor because Wal-Mart is an organization and because "rooting itself" can only *literally* be done by a plant.

In this whitepaper, we justify this approach, describe in some detail how metaphors can be detected and analyzed in this way, and suggest how the same methodology can be extended to support increasingly sophisticated processing of metaphorical language. Finally, we describe how one might begin to validate applicability of automated metaphor understanding for social science research.

2. Technical Challenges and Approach

Metaphor is ubiquitous. Superficially, it decorates language and makes reading and listening a more pleasant experience. Beneath the surface, it provides the basis for many of the most common linguistic constructions. And even more deeply than that, it can and does license reasoning, enabling the reader or listener to make enormous inferential leaps to conclusions which, while not sound, are quite often correct (though sometimes tragically wrong). Analogy can get one to the same destinations, but less directly.

Instead of thinking of metaphor as icing on the cake, it plausible that metaphors build on direct somatic experiences we have as neonates, gradually leading us to higher cognitive processing and the use of language. This notion, perhaps most famously advanced in *Metaphors We Live By* (Lakoff & Johnston 1980), establishes metaphor as a primary source of meaning in language.

There is reason to believe that cultural differences, and even the conceptual structures of individuals, are reflected in our use of metaphor (Yu, 1995, Kövecses 2005), but, so far, automation has not been applied to support a systematic analysis of the relationship between metaphor other objects of interest to social science research.

2.1. Metaphorical analysis can be automated

Ever since computers were conceived as general symbol manipulation systems, there have been forays into automated understanding, production, and use of metaphors (Minsky 1968). It played a key role in establishing and setting the direction for Cyc (Lenat *et al* 1983). The largest recent effort to systematically automate the analysis of metaphor has been carried out by Pasanek and Sculley (2008), who, taking advantage of Pasanek's previous effort to compile a corpus (most recently reported in (Pasanek 2009)) of metaphors of mind, were able to train support vector machines to identify further examples of such metaphors, from lexical bag-of-word representations of sentences. This work has been important: it automated the detection of metaphor in a way that generalized across authors, demonstrating the likely feasibility of richer and more comprehensive forms of automation.

Despite that history of relevant work, there is still much room for improvement and extension and further research and development; the earlier work has generally focused on a single class or a small range of metaphor; and it has generally used a somewhat unsophisticated NLP analysis. The earlier efforts therefore do not solve the problem of metaphor, but they can be harnessed as partial solutions.

2.2. Metaphors can be detected because they break the rules

What is needed, then, is a more universal indicator that metaphor is being used; one that does not need to be trained with manually labeled data for each new metaphor class. Such a feature can be found in semantics: a metaphor can be seen as a particular sort of analogical mapping, signaled by the violation of concrete constraints on relationships between the concepts denoted. In contrast with theories of metaphor that focus on the meanings of words[1], Ricoeur (1975) characterizes metaphors as using "impertinent predication" – the semantics of predicates in sentences are violated to signal, and construct, a metaphor that connotes an intended meaning. To illustrate, consider the corresponding Mandarin and English sentences:

[1] These theories suppose that metaphors are a form of "deviant denomination" with tropes, or remote variations on the meaning of words, being used for decorative purposes.

Zhe-xia tiqi le wo-de xingzhi / **This time it lifted my mood.** (Yu, 1995),

These both refer to a physical process, *lifting an object*, applied to an intangible mental substance, mood. If one takes this literally, what is described is clearly impossible. And, using logical deduction, one can *prove* that it is impossible. Similarly, we have in Hungarian and English:

Túlcsordult a szíve a boldogságtól / **His heart overflowed with joy.** (Kövecses 2005),

a sentence in which, again, a physical process, overflowing, is impossibly ascribed to an intangible thing, mood. Replace *joy* with *blood,* in that sentence, and the violation goes away (it becomes a medical problem rather than a linguistic one.)

It is straightforward for people to prove that these statements, taken literally, cannot be true, but a little less straightforward for computers. The proofs depend on a prior background theory against which to test plausibility. Cyc (Lenat & Guha, 1980; Matuszek et al, 2006) provides one such source of the theories that can be used to detect non-literal usages, and an inference engine that supports the necessary deductions, and will be used in the examples. Other knowledge sources, with similar content and other deduction mechanisms with similar power could be used, if available, or could supplement the reach of Cyc.

To make the process concrete, analysis of the first metaphor, above, proceeds as follows: First the Cyc term "LiftingAnObject",is identified as the denotation of the word "lift" when used as a verb, and the sentential argument structure is identified:

```
(denotation Lift-TheWord Verb 0 LiftingAnObject)

(verbSemTrans Lift-TheWord 0 TransitiveNPFrame
    (and
      (isa :ACTION LiftingAnObject)
      (objectActedOn :ACTION :OBJECT)
      (doneBy :ACTION :SUBJECT)))
```

In the background theory for the *concept* LiftingAnObject the objectActedOn is, not unreasonably, required to be PartiallyTangible, by way of:

```
(interArgIsa1-2 objectActedOn LiftingAnObject PartiallyTangible)
```

Moreover, "mood", which has the denotation, "FeelingAttribute":

```
(denotation Mood-TheWord MassNoun 0 FeelingAttribute),
```

is known to be Intangible. Finally, another assertion in Cyc's knowledge base holds that Intangible and PartiallyTangible objects are disjoint, allowing Cyc to infer (in 56ms) that the sentence cannot be literally true:

```
Mt : InferencePSC
(implies (and (isa ?X FeelingAttribute) (isa ?Y LiftingAnObject))
  (not (objectActedOn ?Y ?X)))

Query was proven True
```

2.3. Processing at Scale: Large-scale Shallow Information Extraction

It is, of course, not possible to applying surface semantic incoherence as a means of detecting metaphor if the predicate/argument structure of text cannot be analyzed. Fortunately, there has been recent work on doing this at scale. TextRunner (Banko & Etzioni, 2008), for example, has been used to extract predicate/argument structure rapidly from the text in 9

million web pages. To illustrate, a search for the predicate "overflow", using a TextRunner API, yielded thousands of examples, including:

> <arg1 toks="the small courtroom" pos="NN" mor="**courtroom**" syn=""/>
> <pred toks="were overflowing with" mor="**be overflow with**" typ="BNP" syn=""/>
> <arg2 toks="spectators" pos="NN" mor="**spectator**" syn=""/>
> <url>http://www.nyu.edu/library/bobst/research/tam/gardner_mss.html</url>
> <just>**He cleaned his glasses and slowly inspected the small courtroom , where seats were overflowing with spectators and newspaper men.<**/just>

The arguments, "courtroom" and "spectator" can be readily tested for coherence with literal "overflowing". While a courtroom is, indeed, provably a container, a spectator, by virtue of being a **PerceptualAgent-Embodied** cannot be a **LiquidTangibleThing**[2] The extracted relation, then, cannot be literally true. It must be a metaphor[3].

Other TextRunner-extracted arguments of "overflow" include (inboxes, spam), (church, leader), (emergency department, patient), (folk, stories), (glass vase, spray chrysanthemums), (pisoirs, drug deal), (hospital, victim). None of these presents any great barrier to checking the requirement for generalization to (container, LiquidTangibleThing). As a group, they support the kinds of analysis envisioned in the RFI: "overflow", when used metaphorically, tends to be coherent in its first argument (they are, in fact, containers) but not in its second (they aren't liquids)[4]. To be clear, literal uses of "overflow" do appear: for example (bath, water), (toilet, water), (gutter, water), and can be shown to support the constraints on the word's literal meaning. In one borderline example, "roads overflowing with water" (roads, water), it is unclear even to these human authors whether the use is literal or not.

Further, suggesting that the English version of the example from Kövecses 2005, above, could easily have been found by such automated analysis, a search of the TextRunner DB finds hearts overflowing with gratitude, love, and gladness; in Cyc, these are all known to be specializations of "FeelingAttribute", and knowing thish is already a useful component of an analysis of the metaphor. It would not be surprising to find further, across cultures or subcultures, that hearts are said to overflow with particular emotions with different frequencies.

2.4. Processing at Depth: Harnessing both Logic and Probability

If the sentence said, instead, that the spectators were *crushed* into the courtroom, a longer inference chain would be needed to conclude that while conceivable, it's probably not literally true, and this raises the important issue of combining symbolic logical reasoning (and discourse analysis) with probabilistic reasoning, to detect and understand metaphors which could – just barely – be literally true but would have been "the lead" of the story if they were nonmetaphorical; e.g., Jay Leno *slaying* his audience.[5]

---

[2] For the curious, this is because a **PerceptualAgent-Embodied** is a sort of **Organism-Whole**, and a **FluidTangibleThing** is a kind of **HomogeneousStructure**, and another Cyc axiom states that never the twain shall meet: (disjointWith HomogeneousStructure Organism-Whole).

[3] Or a processing error in information extraction.

[4] Not always though: e.g. (cellar door, quality wine).

[5] This is certainly not all we have to say on this topic. Indeed, in the long term, our research program depends on the development of effective probabilistic logics and their means of application. Conventional logic allows us to approach the elements problem of metaphor in with techniques that are well developed, and we focus on that approach here. Later work will explore the use of probabilistic reasoning in metaphor processing in more detail.

2.6. A Community Through Time is like a Set of Communities.

Pasanek (2009) demonstrated the feasibility of relatively shallow metaphor analysis, but it also suggests a means of mitigation for a technical barrier to validating our approach for social science applications, in particular for understanding the use of metaphor, and its conceptual implications, across communities.

Such research is made more difficult by the fact that it is not yet straightforward to obtain comprehensive lexical-semantic mappings to a background knowledge base for languages other than English. Multiple language versions of WordNet (Fellbaum, 1998), and FrameNet (Baker et al 2003), and resources such as Wiktionary (2009), offer one possible approach, but the Pasanek and Sculley research program (Pasanek 2009, Pasanek & Sculley 2008) suggests another: the authors aimed identify and to explain shifts in the use of particular metaphors over time. For testing analytical techniques based on semantic metaphor detection, then, it may be possible to use the cultural difference over time displayed in out-of-copyright English language literature, which is being made available by e.g. Project Gutenberg (2009) and Google Books (2009), as a proxy for contemporary cultural differences.

2.7. Restoring Order: Computing Metaphorical Transformations

Detecting that a metaphor may be present is certainly useful. One could, for example, simply count the frequency of terms occurring in implausible relations to get an idea of whether the concepts involved in metaphors vary from culture to culture. Or, by using a semantic network, like WordNet or OpenCyc, one could look for differences in the types of things that are connected by metaphor; in the Hungarian example above, for example, the concept Heart generalizes, *inter alia*, to the Cyc concept InternalAnatomicalPart, and the concept Happiness ('Joy') generalizes, *inter alia*, to FeelingAttribute. Perhaps metaphors relating anatomical parts to feelings are more, or less common in Hungarian culture; if so this simple analysis would uncover the fact.

While this is intriguing, an even more interesting analysis, however, is enabled if we can automate the process of identifying which *class* of metaphor is in play. For example, detecting that a passage instantiates the Lakoff & Johnson (1980) metaphor "SOCIAL ORGANIZATIONS ARE PLANTS" or "LIFE IS A JOURNEY" would enable a statistical analysis of the relative rates at which these metaphors occur across cultures[6] and subcultures (time, industry, age, occupation, etc.), and an analysis of *how* they are instantiated ('Which *kinds* of plants?', 'Which *mode* of journey?'). Perhaps most importantly, it might enable the analysis of articulated metaphors, where a metaphor is extended over several locations in a document.

The idea behind the proposed computational metaphorical transformations is to generate virtual objects, in a discourse context, for which the metaphorical statements can be taken literally. This could be done in real time, as needed, or, even better, a large cache of automatically-generated metaphorical transformations could be produced and indexed for efficiency.

Concretely, a metaphorical transformation is a function from instances of one type to instances of another that maintains meta-information about the latter instances allowing them to be connected to their source in the discourse.

---

[6] For an example of such differences in the rate of metaphor use, see Niki Köves, 2002 in Kövecses (2005) Table 4.1, which ranks use of metaphors about life between Americans and Hungarian speakers; Americans favor "LIFE IS A PRECIOUS POSSESSION" first, and Hungarians "LIFE IS A STRUGGLE/WAR" first.

Consider the sentence, "*Standard Bank planned to grow its branch network in Nigeria by 20 percent by the end of the year*."[7]. The "SOCIAL ORGANIZATIONS ARE PLANTS" metaphor is so common that one might (and many dictionaries do) treat the terms "grow" and "branch" as simply being independent word senses applied to organizations and their parts. But that approach does not scale up; at best it keeps almost catching up:— fresh metaphors roll out of the oven steadily. We propose that a better understanding of the text can be had by processing the metaphor directly. A functional logical term, such as MetaphoricalMappingFunction(SocialOrganization, Plant, StandardBankBranchNetwork[8]), can be used to represent the metaphorical plant instance in question, allowing a direct interpretation of the sentence. In Cyc, this involves applying (verbSemTrans-Canonical Grow-TheWord 0 TransitiveNPFrame HusbandryOfPlant (TheList performedBy objectActedOn)) to interpret the subject, Standard Bank, as the performer of plant husbandry, and the metaphorical plant as the objectActedOn.

For ambitious applications of NLP, this sort of mapping should serve as a powerful means of supporting textual inference for full text interpretation: if there are known to be consequences, to a plant, of being acted on in HusbandryOfPlant, then these consequences also apply to the metaphorical plant. Since the connection between the original organization and the metaphorical plant is kept, inverse transformations can be used to make inferences about the state of the branch network after growth: for example that it has more sub-organizations or members[9].

For simpler applications, this sort of transformation can be used more shallowly. Consider the the following conceit: "Once Starbucks took root outside Seattle, it grew rapidly throughout the US, sprouting in the fertile ground of sophistication-craving suburbs". Once the transformation of Starbucks, the organization, into a metaphorical plant has been forced by the first clause, it becomes straightforward to assign it a role in the growing, and sprouting events: they require a plant, and to establish its relationship to the (also metaphorical) fertile ground. This role assignment permits subsequent analysis of the types of organizations that appear in the various instantiations of the "ORGANIZATION IS PLANT" metaphor (growing, sprouting in fertile ground, taking root) and of the ways those aspects are articulated. One might anticipate, for example, that culturally agrarian societies would use more elaborately articulated horticultural metaphors, based on a more sophisticated background theory.

2.9. Notes on Evaluation

The most obvious evaluation criteria are nevertheless quite worth applying: Perhaps the most straight forward (1) is counting the frequencies of correct and incorrect identification of metaphors in text (e.g., by hand-scoring the systems' decisions on a sample of the output of TextRunner). Two more sophisticated criterion would test whether a system (2) is producing the correct answers in (1) for the right reasons, citing metaphorical transformations which coincide with the human evaluators' independent introspection, and (3) score a system's

---

[7] From South African online news outlet "BusinessReport", accessed Sept 8th 2009 at http://www.busrep.co.za/index.php?fSectionId=563&fArticleId=5155431

[8] Identifying this concept and its relationship to the verb "grow" is within the reach of current NLP – as a reader can verify at http://www.link.cs.cmu.edu/link/submit-sentence-4.html by parsing the sentence.

[9] Mapping from the consequences of inferences (a theory) about the metaphorical object back onto elements of a theory about the real thing it represents is, perhaps, the central problem in applying this notion of computational metaphorical transformation to deep understanding of connected text. This topic will be developed further in subsequent papers.

ability to classify metaphors appropriately (e.g., into "INANIMATE AS AGENT" or specializations thereof such as "COUNTRY AS PERSON"). Finally, one could measure the system's ability to posit new metaphor classes, but clustering forms of Ricoeur's impertinent predication. Together, these four evaluations – which can be performed at modest cost by leveraging already-built technologies such as Cyc, TextRunner, WordNet, and FrameNet – would establish that the metaphorical analysis is sufficiently deep to serve as a basis for applications in deep language understanding.

3. References:

Baker, C. F., Fillmore, C.J. and Cronin, B. (2003) The Structure of the Framenet Database, *International Journal of Lexicography*, Volume 16.3: 281-296.

Banko, M., and Etzioni. O. (2008). The Tradeoffs Between Open and Traditional Relation Extraction, In *Proceedings of the 2008 Annual Meeting of the Association for Computational Linguistics* (ACL), Columbus.

Fellbaum , C, (Ed) (1998) WordNet: An Electronic Lexical Database, MIT Press

Google Books (2009) http://books.google.com/googlebooks/agreement/ accessed Sept 8th 2009.

Gutenberg (2009) http://www.gutenberg.org/wiki/Main_Page accessed Sept 8th 2009

Kövecses, Z., (2005), *Metaphor in Culture: Universality and Variation*, Cambridge University Press.

Lakoff, G., and Johnston, M. (1980), *Metaphors we Live By,* Chicago, University of Chicago Press

Lenat, D.B., Borning, A., McDonald, D., Taylor, C., and Weyer, S. (1983) Knoesphere. *Proc. IJCAI-1983*. 167-169

Lenat, D.B. and Guha, R.V. (1980) *Building Large Knowledge Based Systems*, Reading, Massachusetts: Addison Wesley, 1990.

Matuszek, C., Cabral,J., Witbrock, M., DeOliveira. J., (2006) An Introduction to the Syntax and Content of Cyc In Proceedings of the 2006 AAAI Spring Symposium on Formalizing and Compiling Background Knowledge and Its Applications to Knowledge Representation and Question Answering, Stanford, CA, March 2006.

Minsky, M. (Ed.), (1968) *Semantic information processing*, Cambridge, MA: MIT Press.

Pasanek, B., and Sculley, D. (2008) Mining millions of metaphors, *Literary and Linguistic Computing* 2008 23(3):345-360; doi:10.1093/llc/fqn010

Pasanek, B. (2009) *The Mind is a Metaphor Database,* accessed 2009/09/07 http://mind.textdriven.com/db/browse.php

Ricoeur, Paul. (1975). *The Rule of Metaphor: Multi-Disciplinary Studies in the Creation of Meaning in Language*, trans. Robert Czerny with Kathleen McLaughlin and John Costello, S. J., London: Routledge and Kegan Paul 1978

Wiktionary (2009) http://www.wiktionary.org/ accessed Sept 9th 2009.

Yu, N. (1995) Metaphorical expressions of anger and happiness in English and Chinese. *Metaphor and Symbolic Activity* 10(2), 59–92.