12-1-2008

Abstraction as a basis for the computational interpretation of creative cross-modal metaphor

Sylvia Weber Russell

University of New Hampshire, Durham, swr@cs.unh.edu

Comments

This is an Accepted Manuscript. The final publication is available at Springer via https://dx.doi.org/10.1007/s10772-009-9042-8

Recommended Citation

Sylvia Weber Russell - Abstraction as a Basis for Computational Interpretation of Creative Cross-Modal Metaphor. International Journal of Speech Technology, 11 (3), pp. 125-134 (2009). DOI: 10.1007/s10772-009-9042-8

Abstraction As a Basis for the Computational Interpretation of Creative Cross-Modal Metaphor ⋆

Sylvia Weber Russell

Department of Computer Science
University of New Hampshire
Durham, NH 03824, USA
swr@cs.unh.edu

Abstract. Various approaches to computational metaphor interpretation are based on pre-existing similarities between source and target domains and/or are based on metaphors already observed to be prevalent in the language. This paper addresses *similarity-creating* cross-modal metaphoric expressions. It is shown how the "abstract concept as object" (or reification) metaphor plays a central role in a large class of metaphoric extensions. The described approach depends on the *imposition* of abstract ontological components, which represent source concepts, onto target concepts. The challenge of such a system is to represent both denotative and connotative components which are extensible, together with a framework of general domains between which such extensions can conceivably occur. An existing ontology of this kind, consistent with some mathematic concepts and widely held linguistic notions, is outlined. It is suggested that the use of such an abstract representation system is well adapted to the interpretation of both conventional and unconventional metaphor that is similarity-creating.

1 Introduction

The last couple of decades have seen an increasing number of computational approaches to processing metaphor. By interdisciplinary consensus, this research has generally been implemented as processes that map an expression in a *source domain* (domain of the metaphorically used concept) to an interpretation in a topic- or *target domain* (domain of intended meaning). (Within-) physical-domain metaphor, as in the war horse example *The ship plowed [through] the sea/waves*, received attention early on [Russell, 1976]. Treatments that focus solely on physical-domain metaphor include those by Wilks [Wilks, 1977] [Wilks, 1978] and Fass and Wilks [Fass and Wilks, 1983] and are discussed in [Russell, 1986]. Fass [Fass, 1997], while presenting an extensive treatment of metaphor in the context of literal, metonymic and anomalous expressions, also focuses mainly on physical-domain metaphor.

This paper argues for the role of an abstract ontology in the interpretation of *cross-modal* metaphor, with special attention to *similarity-creating* metaphor. Cross-modal metaphor extends across "conceptual domains" (modes, levels), as in *to quarry ideas*, which involves extension from a physical to a mental domain, in *Encyclopedias are gold mines*, which involves extension from a control (possessive) domain to a mental domain, or *the cadenza of his gymnastic performance*, which involves extension from a sensory (sound) domain to a physical domain.

The discussion begins with an indication of what is meant by similarity-creating metaphor and how some of the major research on metaphor does not address it. Some notes on ontologies follow, as well as observations on abstraction, including mathematical abstraction and its potential relevance to an abstract natural-language ontology. A program that relies on an abstract ontology to address similarity-creating metaphor is then outlined, with an explanation of components of the ontology.

⋆ An earlier version of this paper was presented at the 5th International Workshop on Natural Language Processing and Cognitive Science - NLPCS 2008

This description is followed by brief illustrations of relevant aspects of interpretations of verbal and nominal metaphor. The paper concludes by noting that some other researchers have found it necessary to extend their metaphor processing systems with (at least implicit) abstract ontological components. This trend, along with interdisciplinary observations on abstraction, suggests the need for attention to an abstract ontology in a metaphor processing system, at least if the system is claimed to be explanatory.

2 Similarity-Creating Metaphor

With respect to metaphor, the word 'creative' is often used interchangeably with the word 'novel.' By 'novel metaphor,' some researchers refer simply to metaphors that their systems—or perhaps even humans—have not previously encountered. Such metaphors may be based on representations that capture similarities existing prior to use of the metaphor. By contrast, Indurkhya [Indurkhya, 1992] presents evidence of the significant role of similarity-*creating* metaphors in cognition.

As Indurkhya points out, in many metaphoric expressions where verbs are used unconventionally—some would say in a novel way, there is a pre-existing similarity. For example, in *The sky is crying*, there is an easily recognizable similarity between crying (tears falling) and one of the few things that fall from the sky (rain drops); the metaphor can be analyzed by *comparison*, though it also *suggests* sadness. Similarity-creating metaphor, on the other hand, is characterized as *change of representation*: "In instantiating the source concept network in the target realm, parts of the realm are 'grouped' together and made to correspond to the concepts of the concept network. In this process, the target realm is given a new ontology, and its structure, as seen from the more abstract concept network layer, is changed" (p. 254).

Indurkhya acknowledges that the difference between suggestive similarity-based metaphor and similarity-creating metaphor may be a matter of degree, i.e., degree of participation of the target and source domains. That is, the closer the metaphor is to being similarity-creating, the more the source ontology is imposed and the less the pre-existing target ontology remains. In this paper it is assumed that cross-modal metaphor (which Indurkhya does not focus on) is in a sense always similarity-creating, because the real-world details of source and target will always differ. For example, in the metaphor, *Encyclopedias are gold mines*, there is little literal similarity between encyclopedias and gold mines, or between reading and mining.

3 Approaches to Metaphor

Through metaphor, different source concepts may be used to structure the target in different ways. Lakoff and Johnson [Lakoff and Johnson, 1980] recognize both the "conceptual metaphors" (or "metaphor themes") *LIFE IS A JOURNEY* and *LIFE IS A GAME*, and perhaps other "life" metaphors. Thus similarity can be created by re-conceptualizations. While such conceptual metaphors are certainly of theoretical interest, a problem for computational purposes is that they are categories, without specifications of which components of the source domain of a word are extensible.

The early approach of Carbonell [Carbonell, 1980] [Carbonell, 1982] is based on the stored conceptual metaphors of Lakoff and Johnson. However, systems that rely only on stored conceptual metaphors cannot interpret linguistic metaphors that do not fit any stored conceptual metaphor. (To interpret *The discussion scintillated* do we need a *DISCUSSIONS ARE STARS* or perhaps a *DISCUSSIONS ARE SPARK PRODUCERS* conceptual metaphor?) Also, the metaphoric nature of the transferred properties themselves is not addressed. For example, the phrase 'firmly supported,' used in Carbonell's example of the *MORE IS UP* conceptual metaphor, is simply applied to both source and target domains without semantic analysis.

Hobbs [Hobbs, 1992] addresses metaphor without recourse to stored metaphors, using inferences to express linguistic relationships. In his illustration he matches 'send (a bill)' in a Congress schema to 'pitch (a ball)' in a baseball schema, and "proves" the correspondences between roles in the two schemata. This metaphor is certainly novel, but Hobbs's interpretation process is based on existing similarities between the source and target schemata. His system apparently does not provide metaphoric interpretations in cases where no appropriate schemata exist. The verb 'pitch' is defined as 'send' plus certain unnamed differentiating conditions. It is not clear whether 'send' is a basic representation component or one of many literal and metaphoric senses.

Approaches of other researchers that show some potential to address similarity-creating metaphors are discussed at the end of this paper.

4 Ontologies

The term "abstract ontology" might be seen as an oxymoron, and it is, if an ontology is that which purports to describe the "real world." Wilks [Wilks, 2007], in dismissing the distinction between traditional/classical and modern/AI-type meanings of "ontology" as unimportant for AI/NLP purposes, also rejects any claims that "cleaning up" given ontologies will result in any notable advances in the field. This view (which I accept) is mentioned in order to emphasize that the focus in this paper is only on the role that abstracted components can play in a computational metaphor interpretation system with attention to presumed cognitive components, and on what *types* of components are needed and are peculiar to metaphor interpretation. While the ontology is explained below, the intent is not to justify the exact form the individual components take. It is important, however, that the ontology, being abstract, be relatively small and transparent, for purposes of evaluation and revision.

A cross-modal metaphor-relevant ontology is based not on any objective reality, but on a certain unconventional *view* of reality through language, which itself represents a conceptualization. A perhaps noncontroversial observation of Quine [Quine, 1969] on the ontology of language would seem to apply to abstractions *from* language (i.e., to an abstract ontology) as well—namely, that differences between one person's ontology and another may depend simply on how the ontologies are "sliced" or how components are grouped; correspondences between ontologies will probably not be one-to-one (cf. also Whorf [Whorf, 1956]. There is no claim in this paper, then, that the components of the abstract ontology are universal, uniquely "correct," or language-free; there is merely an appeal to some parallelism to other disciplines and to a consensus of "reasonableness" by speakers of the same language and others that are related to some extent. Neither is there speculation on the source of the given ontology in developmental or evolutionary terms.

In cross-modal metaphor, any perceived or imposed similarity as mediated by the ontology is *abstract* (in the conventional rather than strict use of that word); some considerations of abstraction follow.

5 Abstraction

In abstraction, one representation is converted to another representation in which some details of the source representation are dropped, but the underlying structure, or part of the structure, is left unaltered. In a paraphrase of linguistic metaphor, the abstracted representation can be thought of as an interlingua between the metaphor and its (more) literal translation. In a sense, any representation, whether a mapping between real-world concepts and symbols or between those symbols and higher-level symbols, is abstract. In the context of mathematics learning, Kaput [Kaput, 1989] defines four interacting types of representation - 1) cognitive and perceptual, 2) explanatory representation involving models, 3) representation within mathematics and 4) external symbolic representation, such as a chip, which can be instantiated by many different objects and can thus be a *generalization* or *abstraction* for cookies, baseball cards, dollars, etc. In natural language, similarly, the concept underlying the word 'object' can be thought of as a generalization or abstraction for the mentioned items; it is plausible that the cognitive components which relate to mathematical abstraction are (or overlap with) those which structure linguistic metaphor. As it is being argued for an abstract ontology for metaphor, a consideration of relationships between mathematical and abstract linguistic components that might be included in such an ontology follow.

5.1 Mathematical Language

Mathematics education literature frequently refers to the power of mathematics to account for many analogous situations through its abstract language. It is often difficult to characterize mathematical language and natural language independently in discourse, since mathematical *concepts* can be *embedded* in natural language, not only in mathematical word problems, but in our everyday language about situations. [1] For example, English can embed both explicitly numerical references, such as 'ten' and 'a dozen,' and expressions that are mathematically relevant but not necessarily so intended, such as 'the rest of them,' 'a slice of pizza,' 'altogether,' 'join,' 'more than,' etc. [LeBlanc and Weber-Russell, 1996]. The meshing of these languages corresponds to Kaput's interaction between cognitive/perceptual and mathematical representations and suggests common ontological components. For example, an abstract PART concept can be realized in both mathematical and nonmathematical language.

5.2 Reification as a Basis for Spatial Structuring

In mathematical language, arithmetic equations represent structures with *numbers* as abstractions not only of sets of objects, but also of non-object concepts (as in 'he fell *twice*'), and with *operators* that relate these sets; the abstraction to numbers establishes the basis for the equation. Similarly, in linguistic metaphoric extensions from the physical domain, nonphysical concepts may become abstract "objects," allowing verbal concepts to "operate on" them. Reification (or "nominalization")—treating an action, relation or attribute as an "abstract object" in the form of a noun—is thus a first step in the creation of this kind of metaphor. Expressed syntactically, reification is an instance of the "abstract concept as object" metaphor [Russell, 1989], which has become integrated into (some) natural languages as dead (frozen, assimilated) metaphors, i.e., metaphoric language usually thought of as literal. Mathematical language and much of metaphoric language can thus be viewed as sharing spatial grounding,[2] suggesting that not only physical but also nonphysical verbal concepts might be analyzed in terms of spatial structures.

To illustrate, the physical action underlying the verb construct 'chase out/away (e.g., mosquitoes)' can be extended to apply to conceptually different types of objects. In 'chase away an idea,' the 'idea,' which is a reification in a mental domain, is "taken away" from the thinkers of the idea; mathematically, to 'chase away six mosquitoes,' as in a word problem, may mean to subtract or "take away" 6. In both cases, symbols are mapped from one domain to another through an abstraction representing "leaving a state" (of thought or of the presence of the six mosquitoes).

Prepositions can also be interpreted metaphorically as well as mathematically if the head of the prepositional phrase is reified. Figure 1 shows how the preposition 'of' can be understood linguistically as a metaphoric 'part' or quantitatively as 'subtracted from.' The extension of 'of' in combination with syntactic reification or quantitative abstraction (middle level of the ellipses) provides the basis for metaphoric and mathematical expressions respectively as shown.[3]

Reification, then, enables the natural-language extension similarly as quantification enables the mathematical extension.

Thus if we settle on a basic set of abstract components in terms of certain state configurations and use them as the basis of abstract verb definitions applying to multiple domains, then these

[1] See Cummins et al. [Cummins et al., 1988] for evidence of young children's confusion about such embedded language.

[2] See Section 8 for an alternative to this assumption.

[3] See Brugman and Lakoff [Brugman and Lakoff, 1988] for a systematic treatment of extended senses of prepositions.

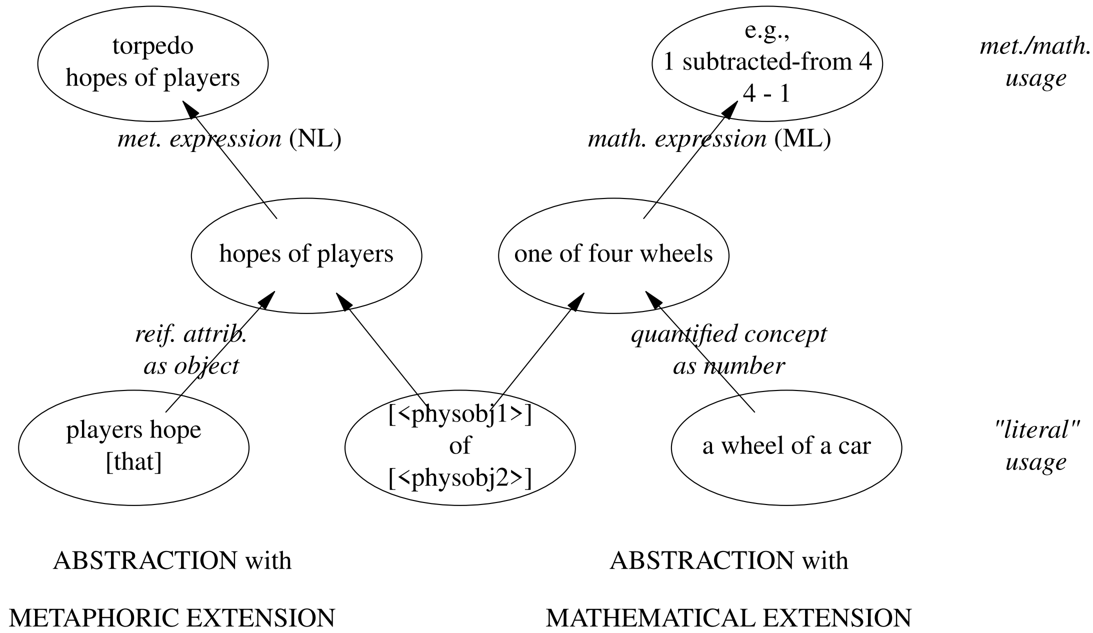

Fig. 1. Metaphoric and mathematical extensions of the preposition'of.'

components can be considered to be extended in metaphor and to contribute to its interpretation (see Section 8).

6 An Ontology-Based Metaphor Analysis Program

MAP, a computational metaphor paraphrase program [Russell, 1976] [Russell, 1986] [Russell, 1992], interprets a simple, partially metaphoric sentence in terms of a roughly equivalent paraphrase conventionally considered as "literal." The most critical aspect of the program resides in the (abstract) lexicon, where verbs and nominals are represented by components of an abstract ontology. For verb-based metaphor, components representing a verb which serves as a metaphoric *source* concept are interpreted in the *target* domain as indicated by the nominal concepts with which the verbal concept is used. Thus for *She chased away the thought*, a mental domain is indicated by 'thought,' and the primitives underlying 'chased away' lead to a paraphrase including a phrase such as 'voluntarily stopped thinking about.' For nominal metaphor, the primitives underlying *salient properties or predicates* [Russell, 1986] of the source nominal are transferred to the target representation. Thus for *Political movements are glaciers*, the potentially extended predicates include components representing slow change.

MAP treats dead metaphors and novel metaphors (whether similarity-based or similarity-creating) in the same way, though of course dead metaphors and even some metaphors that are "alive" but conventional could be defined directly in the lexicon for efficiency purposes. The focus of this discussion, however, is on MAP's ability to interpret similarity-creating metaphors.

The assumption that natural language and mathematics share spatial structure, e.g., structure in terms of objects and relations, suggests that 1) a small number of abstract descriptors that overlap with those of mathematics in being spatially based reflect some intuitive consensus of speakers of the language with respect to the design of an ontology and 2) such spatially based structures provide a framework for additional, qualifying primitives, some of which also correspond with mathematical concepts. The ontology described below consists of *extensible components* including spatial structures, which represent the potential similarities between source and target, and *domains*, which represent the differences.

6.1 Abstract Extensible components

The task of determining a set of extensible components of verbal concepts entails considering which concepts speakers of a given language recognize in a literal meaning of a verb that allows them to understand a metaphoric use of that verb, even if they have never heard it before. If much of our language is spatially structured, we should be able to suspect some cognitive basis for components in the abstract domain of mathematics (arithmetic, calculus, logic). Elementary physics concepts applicable to a wide range of everyday physical situations can also be expected to play a representational role. The following structures and features either have a math-physical counterpart and/or have a broad linguistic consensus.

Structures All verb structures are based on a STATE. The STATE may take the form of either the existence of a nominal (OBJECT BE); the relation of a nominal in relation to other concepts (OBJECT AT LOCATION; or a (static or dynamic) attribute of a nominal (OBJECT BE <attribute>).[4] These abstract STATE structures might be thought of logically as one- or two-argument predicates or linguistically as unary or binary abstract case structures.[5] In addition to negation (NOT), components may represent the beginning or end of a STATE, or transition through a space, which in mathematical functions correspond to boundary points (limits) or a path between two points respectively. Any STATE structure can be an effect or *result caused by* an event or by an (animate) AGENT. It is the result portion of the abstract structure of a verb involving causation that is considered of primary salience and receives a domain specification. Verbs requiring representation including other higher-level primitives, such as purpose, conjunction, temporal sequence, etc. have not yet been included in MAP's lexicon.

Features It is qualifiers and connotations that are often the point of a metaphor. These are represented as abstract, "conceptual" features—more flexible than explicit categories and perhaps accounting for some of the fuzziness perceived in the concepts they apply to. These have polarity or magnitude specifications as appropriate. As qualifiers of actions, features for action verbs correlate with some mathematical descriptors: CONTINUITY, REPETITION (frequency) and SPEED (rate). Verbs with quantitative attributes (e.g., 'grow'), may have MAGNITUDE and GREATER/LESS-than. VOLITION is a feature describing an actor. Responses of an experiencer (reader or hearer) of the metaphor have EVALUATION values (POSITIVE, NEGATIVE) and FORCE magnitude (HIGH, LOW). EVALUATION and FORCE correspond to Osgood's [Osgood, 1980] evaluative and potency factors—two of the three nonstructural factors (the other being activity, refined in the action features above) he empirically determined to be extended in metaphoric usage (see also Aarts and Calbert [Aarts and Calbert, 1979]). Various emotions are also incorporated. Emotional states as real-world concepts are not "abstract" in the sense that NOT or MAGNITUDE is. However, they are clearly extensible, though with varying intensity (the fear experienced when one's hope is torpedoed may not have the same intensity as that when one's boat is literally torpedoed).

[4] Attributes, such as 'red,' 'asleep' or 'hopeful,' "fill in" specific properties of the world and are not necessarily considered "primitive."

[5] Abstract case structures are simpler than traditional case structures, since Fillmore's [Fillmore, 1968] dative and locative cases, for example, or Schank's [Schank, 1975] "conceptual" Recipient and Direction cases are combined in (abstract) LOCATION.

6.2 Conceptual domains

Conceptual domains are orthogonal to the extensible portion of the ontology. For cross-modal metaphor, the specified domains are only those general, Aristotelian-like domains which, along with the PHYSICAL (animate and inanimate), are thought of as human faculties: MENTAL, with subdomains *intellect, attitude* and *will*; SENSORY, with sense-specific subdomains; and CONTROL, with subdomains *intrinsic* (e.g., 'talent') and *extrinsic* (either control of physical concepts or control of action possibilities dependent on others, such as 'rights,' 'duties,' etc.). The same categories are used for nominals, with the addition of TIME and SPACE. This taxonomy within the ontology is obviously breadth- rather than depth-oriented.

Every verb in the lexicon is assigned the conceptual domain in which it is thought to be literal. The model allows a concept in any conceptual domain to be a source, though the source is more often PHYSICAL. It is the difference between the conceptual domains of a verb and the OBJECT of the structure that triggers cross-modal metaphor recognition.[6]

The small size and transparent organization of the set of extensible and nonextensible components allows the management of the ontology and the task of representation in terms of that ontology to be feasible. Also, by defining words through the abstract components, we can note which components, when imposed on the target domain, positively or negatively affect the interpretation of similarity-creating metaphors.

7 Interpretation

7.1 Constraints

MAP does not *compare* a source representation with a target; it is not similarity-based. Rather, the abstract source representation is *imposed*, i.e., directly projected onto the target. However, source representations cannot be imposed arbitrarily. There are some coherence constraints on interpretations to assure (as far as possible) that the expression makes sense metaphorically, i.e., is not "anomalous," indicating a probable mis-parse of the text containing the sequence being analyzed. For example, when a transitive verb is used metaphorically with an object nominal in a different conceptual domain, there are some abstract constraints (analogous to "selectional restrictions" on nonmetaphoric language) that the object must satisfy. These constraints are realized in MAP as abstract conceptual features of nominals. For cross-modal metaphor interpretation, these features are fewer than literal semantic features of nominals, since many details of the nominal concept drop out of the picture. For example, PART (of), CONTAINED (in) and FIXED (to) features merge with PART (of) in nonspatial language, since certain topographical features of spatial objects do not apply. MAP's feature set and its application are discussed in detail in [Russell, 1992]. Current (binary-valued) features are:

- SHAPE (discrete vs. amorphous)
- 1-DIMENSIONAL (linear-like)
- PART
- COMPLEX
- FLUID
- ANIMATE (dynamic).

In addition, the FUNCTION of an artifact (or TYPICAL ACTION of a natural concept) is specified in terms of an abstract verbal representation. While a specific function of an object represents world knowledge, the *concept* of FUNCTION is basic to natural language processing in that it expresses the commonly implied relationship between the object and its user.

[6] That the meaning of a novel metaphor depends on its literal meaning does not necessarily imply that literal meanings are always accessed *before* metaphoric ones by humans.

7.2 Paraphrase

If there happens to be a verb in the target domain that has an abstract representation in common with the source or part of the source (at least the structure), then that verb can be included in the paraphrase. For the example *news torpedo his hope*,[7] given that the lexicon has a sufficiently refined taxonomy of emotions, that verb might be 'disappoint,' which has the same structure as the verb 'torpedo,' i.e., AGENT (or event) cause (the beginning of) NOT STATE (OBJECT BE), where the OBJECT ('hope') is in the MENTAL-ATTITUDE domain. Remaining components (here, this would be FORCE: HIGH, SPEED: HIGH, EVALUATION: NEGATIVE) would be lexicalized directly. A more reliable though less interesting interpretation is given by direct lexicalization of all components. In an "undoing" of reification in the current example, the reified abstract OBJECT 'hope' from the input is mapped to the *verbal* 'hope' (or the adjectival '(be) hopeful') as part of the paraphrase and the remaining components lexicalized. Three possible paraphrases, depending on structure interpretation, are then:

from structure:

news disappoint him OR
news cause he stop hope OR
news cause he start he not be hopeful

from character of the action:

forcefully, suddenly, negative

Abstraction necessarily entails a loss of information, and the paraphrases produced often seem inadequate in being too general, though "literal" and not wrong. The paraphrases appear, however, to be close in content to the types of responses Gentner and France [Gentner and France, 1988] observed empirically, which they classify as "minimal subtraction" (of meaning from the verbal concept in its original sense). In any case, it was deemed important to start with a broad, non-*ad hoc* framework, rather than to attend to target-domain detail.

The characterization of nominals for nominal metaphor interpretation is much more open than for verbs, since objects can mean many things to many people. As the meaning of even one sense of a nominal is less constrained than that of a verb, which has inherent structure, there are more possibilities for similarity-creating metaphors. For nominal metaphor, MAP transfers putative salient properties of source nominals [Winston, 1978] [Russell, 1986], to the target. One of the most prominent properties of a nominal that enters into metaphoric interpretation is its FUNCTION (cf. Gibson's "affordances" [Gibson, 1977]) or TYPICAL ACTION as described above. This property, like any verbal expression, can be represented in terms of the described ontological components.

As nominal metaphor typically involves extension of underlying verbal or attributive components, a brief indication of nominal metaphor interpretation will serve to illustrate further representational aspects of verbal metaphor as well. For the example *Encyclopedias are gold mines*: The abstracted FUNCTION of 'gold mine' (one takes gold out of it) is defined in the lexicon with the resultant STATE structure and connotation:

STATE: (OBJECT:(CONTROL of)gold AT LOCATION:<user>)
EVALUATION: HIGH

The conceptual domain of this FUNCTION, since possession and not simply location is involved, is CONTROL, consistent with the OBJECT as shown. The abstract structure *and* the EVALUATION are integrated into the FUNCTION predicate of 'encyclopedia,' (one reads it, i.e., takes information from it), giving the new resultant STATE structure and connotation

STATE: (OBJECT:MENTAL-INTELLECTUAL AT LOCATION:<user>)
EVALUATION: HIGH

[7] Irrelevant grammar-related elements are ignored in input and output examples.

where the OBJECT is specified only by its conceptual domain, MENTAL-INTELLECTUAL, which is the domain of the target STATE itself. The entire paraphrase then is 'One read encyclopedia has result one has knowledge which-has high value.'

This example is of only average richness, but the added connotation of high value, along with the lack of pre-existing literal similarity, makes it similarity-creating. A metaphor that is perhaps more clearly similarity-creating is the metaphor *Dumps are gold mines*. Here the entire FUNCTION structure of 'dump' (to put things *into* it rather than literally *or* metaphorically take them out) is overridden; the interpretation is that something of extreme value can be found in dumps. Of note is that in both cases the property of HIGH value, along with other factors such as connotations, are culturally based and constitute the kind of information that Indurkhya [Indurkhya, 1992] claims must be represented in meanings of objects if similarity-creating metaphors are to be interpreted computationally. Nominal metaphor interpretations are considered to be only *likely*, not definitive, since the writer/speaker may have wished to highlight some less obvious aspect of the source. However, metaphors that have more obscure interpretations usually require further elaboration, requiring multi-sentence analysis.

This approach appears to correspond with Indurkhya's view of similarity-creating metaphor; the source ontology is imposed on, as opposed to compared with, the target domain. Moreover, cultural and experiential factors—the imagined experience which Indurkhya claims as missing from computational treatments of metaphor—are included in the imposed concepts derived from the abstract lexicon.

Following are sample paraphrases by MAP, showing both its ability to produce minimal interpretations and what refinements—if possible—need to be made.

sentence: he plow-through elizas proposal
interpretation: HE CONTINUOUSLY WITH-EFFORT READ ELIZAS PROPOSAL

sentence: he torpedo elizas proposal
interpretation: HE SUDDENLY CAUSE FORCEFULLY STOP ELIZAS PROPOSAL BE

sentence: eliza decorate idea
interpretation: ELIZA CAUSE START IDEA BE COMPLEX COMMA WITH POSITIVE CONNOTATION

sentence: country leap-to prosperity
interpretation: COUNTRY SUDDENLY START BE PROSPEROUS

sentence: he inhale idea
interpretation: HE START BELIEVE IDEA

sentence: he inhale discussion
interpretation: HE START PARTICIPATE-IN DISCUSSION

Some observations concern points that are minor but are relevant to representation and interpretation. For example, 'with-effort' and 'forcefully' both derive from a FORCE feature, but the former applies (in the source representation) to the actor of the sentence, while the latter applies to the "recipient" of action, i.e., an affected being. This type of difference also determines the placement in the output of adverbs derived from conceptual features; features applying to the actor are placed soon after the occurrence of the actor, after any CAUSE element in the paraphrase. The character of "abstract" OBJECTs, which are not all of the same conceptual category, and whether an actor causes a change in vs. interacts with an object, are also examples of necessary grammatical housekeeping that make a difference in the output phrasing.

8 Comparative Evaluation

While most computational approaches to metaphor do not address imilarity-creating metaphors, as they are not based on a semantic analysis that allows ontological components to shape the target domain, the following research has some aspects corresponding to aspects of MAP.

Martin's [Martin, 1990] system is similar to Carbonell's earlier work (see Section 3), with a more comprehensive implementation. However, he has extended his system through a recognition of the conceptual relationship between states and their beginnings and endings. These correspond to MAP's basic abstract structural components.

Carbonell and Minton [Carbonell and Minton, 1983] specify their method for metaphor interpretation in terms of transfer of portions of a graph consisting of concepts linked by relations. Thus for *X is a puppet of Y*, the CONTROL relation between the object 'puppet' and the actor 'puppeteer' is transferred to the node between X and Y. This process and type of representation is similar to that of MAP. However, a comprehensive representation system does not appear to exist, and they do not incorporate affective or cultural components.

The idea underlying the verb representations of the system of Suwa and Motoda [Suwa and Motoda, 1991] is perhaps the most similar to that of the verb representations of MAP. Their ontology itself does not explicitly distinguish domains as in MAP and thus is not as transparent as MAP's, but they do use a finite, relatively small ontology consisting of what they call abstract primitives. These are only in the form of verb *structures*, through which they *match* source and target verbs—a method which apparently succeeds in an interpretation only if such a match exists. Experiential factors are not incorporated. As it stands, then, their system does not address similarity-creating metaphor; however, they discuss the addition of new components to the target and could in theory achieve this, given their abstract ontology.

In the recent work of Barnden et al. [Barnden et al., 2003] [Barnden et al., 2004] and Agerri et. al. [Agerri et al., 2007], it is acknowledged that many metaphoric usages are not adequately covered by Lakoff's conceptual metaphors. They present "view-neutral mapping adjuncts" (VNMAs), which "transfer those aspects that are not part of any specific metaphorical view" or conceptual metaphor [Agerri et al., 2007]. VNMAs appear to correlate with the structural metaphoric extensions of MAP, and are applied as "default rules."

The metaphor theory and attendant hypotheses underlying the system of Narayanan [Narayanan, 1999] have significant similarities with MAP, though his model differs in his neural-like implementation. As aspects of his theory in part apply to Barnden et al.'s and Agerri et al.'s work as well, it will be discussed in somewhat greater detail. Narayanan's treatment of nominals, verbs and adverbs in verbal metaphor in terms of invariant components corresponds with that of MAP in at least two ways. First, the prevalence of spatio-temporal structures as extensible to other domains is incorporated. (Narayanan proceeds further to establish correspondences between motion verbs expressing such structures and possibilities as part of a sequence of actions and *inferences leading to a goal* in a target domain.) Second, from looking at databases, Narayanan has concluded the invariance of certain "parameters" which correspond to MAP's adverbial features expressing evaluation, agent attitude/intent and other (nongrammatical) aspects. His determination can be viewed as corroborating support for the inclusion of such features.

The fine granularity of Narayanan's representation of his two target domains, e.g., economic policy and politics is a positive feature, though the scope of his system is limited in breadth. While Narayanan works out specific mappings to his target domains, MAP deals *generally* with metaphoric extension between domains in a proposed domain ontology. MAP thus reveals how the representation of a source domain concept in a sentence might structure a target concept in *any* other domain.

Another difference concerns the way in which source concepts are projected metaphorically onto the target. In Narayanan's system, entities and actions are projected directly through pre-established "conceptual metaphors" in the sense of Lakoff, such as *MOVERS ARE ACTORS* or *OBSTACLES ARE DIFFICULTIES*, which must be stored. From the point of view of language understanding, MAP shows how a metaphoric usage might be understood in terms of perceived or imposed similarities represented by semantic components of literally understood lexical items, whether the metaphor is conventional or creative, and whether stored or not. Apart from these explanatory differences, Narayanan's system for projecting verbal concepts has similarities in concept to MAP, with more detailed paraphrases for the two domains he treats. The differences perhaps reflect the differing intended tasks—narratives within a specific topic domain/discipline in the case of Narayanan's system, and spontaneous references to metaphor in open discourse in the case of MAP.

In addition to a comparison of the various computational treatments of metaphor, it is of interest to consider how MAP's components compare with those evolving from theoretical insights on the relationship between syntax and semantics and between mathematical and nonmathematical concepts. Jackendoff [Jackendoff, 1983] and Bouchard [Bouchard, 1995] differ significantly between themselves with respect to the autonomy of semantics from pragmatics. This question aside, MAP, like Jackendoff and Bouchard, aims for a certain level of "conceptual structure" that attends to putative cognitive constraints as well as linguistic phenomena that go beyond what is usually thought of as literal.

MAP's ontology is consistent with the criterion of Jackendoff that builds on the hypothesis of Gruber [Gruber, 1965] in generalizing across modalities. MAP's OBJECT, LOCATION and structural components in general accord with the key role of this hypothesis in the semantics of motion and location. In MAP, the conceptual domain of a simple sentence is determined by the domain of the nominal(s) it includes. Thus as in Jackendoff's hypothesis, semantic fields differ only in the nature of the OBJECT nominals that appear in the conceptual structures of events and states. This element is also consistent with Bouchard's view, in which different contexts give different meanings to the use of the same verb.

Although the purpose of Bouchard—to unite semantics with grammatical constraints—differs from that of MAP, MAP shows an analogy with his syntax-based model in that (in contrast to Jackendoff's cases) MAP has only one or two case-like entities—OBJECTs and LOCATIONs—in the conceptual structure of a simple sentence. An example of MAP's divergence from Bouchard in the choice of representation components lies in Bouchard's claim that to represent a verbal concept only in terms of its cause and result omits any specification of an action that might be involved. However, while MAP does not represent such actions as structural components, they are represented in terms of salient conceptual features. While there is no claim that MAP's features are complete, features of this *type* may be sufficient for metaphor interpretation, whereas they might not be adequate to represent all the details of literal language.

With respect to the place of metaphor in language, Jackendoff rejects the view that a cross-modal system makes metaphor possible, because metaphor is "artistic"—a characterization that is at best dubious in light of the ubiquity of mundane metaphor, such as 'break the law.' Like Jackendoff, though from a different standpoint, Bouchard claims that his theory, which treats abstract representation as central rather than as a medium of extension, accounts for various semantic fields (domains) without recourse to metaphoric extension. He links his claim to his rejection of the centrality of space, saying that "concepts not directly grounded in experience are not any more abstract than spatial notions like orientation" (p. 198). For example, Bouchard points out that a mental relation can be built on any verb expressing an abstract relation of contact. With the observation left aside that "contact" appears to be at root a spatial concept, MAP's assumption of the spatial as a frequent source of extension does not imply centrality of space as he defines it (i.e., centrality of space as requiring different primitives for different domains). In MAP's lexicon, spatial senses are *defined abstractly* as well; thus abstract primitives are unified across domains. MAP's version of Bouchard's "abstract contact" (OBJECT AT LOCATION) could thus be thought of as applying to (as an alternative to "extended to") other domains, including the MENTAL domain. The principal difference is that Bouchard would derive a spatial sense from an abstract component, whereas MAP as a natural language processor *abstracts* components *from* the literal sense. The focus on abstraction and on the separation of abstract components from domains is essentially similar.

Aside from specific questions and purposes, then, I would agree with Bouchard, in divergence from the semantic representations of most computational metaphor research, that "only semantic representations based on sufficiently abstract notions are able to account for language use [and I would add in particular for metaphoric usage] adequately" (p. 15).

Finally, it is such abstract notions that provide the most interesting connection between the assumptions underlying the given ontology and the relationship of metaphor to mathematics in cognition, as described comprehensively by Lakoff and Nuñez [Lakoff and Nunez, 2000]. The above references to mathematics do not imply that mathematics is somehow objectively "true" or *a priori* with respect to the ontology; in relating abstraction, metaphor and mathematics, I would accept the assertion of Lakoff and Nuñez that mathematics derives from the human cognitive system, though no specific cognitive claims are being made here. Their approach does differ in framing this derivation, claiming that mathematics *comes from* conceptual metaphors, while the present discussion characterizes mathematics and metaphor as parallel instances of metaphor. In both approaches, however, ordinary—especially spatial—ideas are seen as a grounding for metaphor and mathematics.

9 Conclusion

Some metaphor programs other than MAP produce more detailed interpretations as a result of being similarity-based or being restricted to certain domains. MAP on the other hand was designed for scope rather than detail, not oriented to any specific examples or domains. Its focus on extensible components based on the semantics of the metaphorically used concept enables it to at least minimally "understand" similarity-creating metaphors. The described ontology accounts for both similarities (through extensible components) and differences (between conceptual domains) underlying cross-modal metaphor. Extensible components include not only structures but also connotations and stereotypic experience, imposition of which is offered as an example of what Indurkhya calls a re-structuring by projection of the source concept network onto the target realm. It would seem that the computational interpretation of similarity-creating metaphors with cognitive relevance requires either an abstract ontology of the type presented here or some implicit incorporation of its elements.

References

[Aarts and Calbert, 1979] Aarts, J. and Calbert, J. (1979). *Metaphor and Non-metaphor: The semantics of Adjective-Noun Combinations*. Max Niemayer, Tübingen, Germany.

[Agerri et al., 2007] Agerri, R., Barnden, J., Lee, M., and Wallington, A. (2007). Invariant mappings and contexts in a computational approach to metaphor interpretation. In *IJCAI-MRCS*.

[Barnden et al., 2003] Barnden, J., Glasbey, S., Lee, M., and Wallington, A. M. (2003). Domain-transcending mappings in a system for metaphorical reasoning. In *EACL*, pages 57–61.

[Barnden et al., 2004] Barnden, J., Glasbey, S., Lee, M., and Wallington, A. M. (2004). Varieties and directions of inter-domain influence in metaphor. *Metaphor and Symbol*, 19:1–30.

[Bouchard, 1995] Bouchard, D. (1995). *The Semantics of Syntax: A Minimalist Approach to Grammar*. University of Chicago Press, Chicago, IL.

[Brugman and Lakoff, 1988] Brugman, C. and Lakoff, G. (1988). Cognitive typology and lexical networks. In Small, S., Cottrell, G., and Tanenhaus, M., editors, *Lexical Ambiguity Resolution*. Morgan Kaufmann, San Mateo, CA.

[Carbonell, 1980] Carbonell, J. (1980). Metaphor: A key to extensible semantic analysis. In *ACL*, pages 17–21.

[Carbonell, 1982] Carbonell, J. (1982). Metaphor: An inescapable phenomenon in natural-language comprehension. In Lehnert, W. and Ringle, M., editors, *Stragegies for Natural Language Processing*, pages 415–434. Erlbaum, Hillsdale, NJ.

[Carbonell and Minton, 1983] Carbonell, J. and Minton, S. (1983). *Metaphor and Common-Sense Reasoning, Rep. No. CMU-CS-83-110*. Carnegie-Mellon University, Pittsburgh, PA.

[Cummins et al., 1988] Cummins, D., Reusser, K., Kintsch, W., and Weimer, R. (1988). The role of understanding in solving word problems. *Cognitive Psychology*, 20:405–438.

[Fass, 1997] Fass, D. (1997). *Processing Metonymy and Metaphor*. Ablex, Greenwich, CT.

[Fass and Wilks, 1983] Fass, D. and Wilks, Y. (1983). Preference semantics, ill-formedness, and metaphor. *American Journal of Computational Linguistics*, 9:178–187.

[Fillmore, 1968] Fillmore, C. (1968). The case for case. In Bach, E. and Harms, R., editors, *Universals in Linguistic Theory*, pages 1–88. Holt, Rinehart and Winston, New York.

[Gentner and France, 1988] Gentner, D. and France, I. (1988). The verb mutability effect: Studies of the combinatorial semantics of nouns and verbs. In Small, S., Cottrell, G., and Tanenhaus, M., editors, *Lexical Ambiguity Resolution*. Morgan Kaufmann, San Mateo, CA.

[Gibson, 1977] Gibson, J. (1977). The theory of affordances. In Shaw, R. and Bransford, J., editors, *Perceiving, Acting and Knowing*, pages 67–82. Erlbaum, Hillsdale, NJ.

[Gruber, 1965] Gruber, J. (1965). *Studies in Lexical Relations. Doctoral Dissertation, MIT, Cambridge, MA*. Indiana University Linguistics Club, Bloomington, IN.

[Hobbs, 1992] Hobbs, J. (1992). Metaphor and abduction. In Ortony, A., Slack, J., and Stock, O., editors, *Communication from an Artificial Intelligence Perspective: Theoretical and Applied Issues*, pages 35–58. Springer, Berlin.

[Indurkhya, 1992] Indurkhya, B. (1992). *Metaphor and Cognition*. Kluwer, Dordrecht.

[Jackendoff, 1983] Jackendoff, R. (1983). *Semantics and Cognition*. MIT Press, Cambridge, MA.

[Kaput, 1989] Kaput, J. (1989). Representation systems and mathematics. In Janvier, C., editor, *Problems of Representation in the Teaching and Learning of Mathematics*, pages 19–26. Erlbaum, Hillsdale, NJ.

[Lakoff and Johnson, 1980] Lakoff, G. and Johnson, M. (1980). *Metaphors We Live by*. Chicago University Press, Chicago.

[Lakoff and Nunez, 2000] Lakoff, G. and Nunez, R. (2000). *Where Does Mathematics Come from? How the Embodied Mind Brings Mathematics into Being*. Basic Books, New York, NY.

[LeBlanc and Weber-Russell, 1996] LeBlanc, M. and Weber-Russell, S. (1996). A computer model of the role of text integration in the solution of arithmetic word problems. *Cognitive Science*, 20:357–408.

[Martin, 1990] Martin, J. (1990). *A Computational Model of Metaphor Interpretation*. Academic Press, New York.

[Narayanan, 1999] Narayanan, S. (1999). Moving right along: A computational model of metaphoric reasoning about events. In *AAAI*.

[Osgood, 1980] Osgood, C. (1980). The cognitive dynamics of synesthesia and metaphor. In Honeck, R. and Hoffman, R., editors, *Cognition and Figurative Language*, pages 203–238. Erlbaum, Hillsdale, NJ.

[Quine, 1969] Quine, W. (1969). *Ontological Relativity and Other Essays (The John Dewey Essays in Philosophy)*. Columbia University Press, New York.

[Russell, 1976] Russell, S. W. (1976). Computer understanding of metaphorically used verbs. *American Journal of Computational Linguistics*, Microfiche 44.

[Russell, 1986] Russell, S. W. (1986). Information and experience in metaphor: A perspective from computer analysis. *Metaphor and Symbolic Activity*, 1:227–270.

[Russell, 1989] Russell, S. W. (1989). Verbal concepts as abstract structures: The most basic conceptual metaphor? *Metaphor and Symbolic Activity*, 4:55–60.

[Russell, 1992] Russell, S. W. (1992). Metaphoric coherence: Distinguishing verbal metaphor from anomaly. *Computational Intelligence*, 8:553–574.

[Schank, 1975] Schank, R. (1975). *Conceptual Information Processing*. North Holland, Amsterdam.

[Suwa and Motoda, 1991] Suwa, M. and Motoda, H. (1991). Learning metaphorical relationships between concepts based on semantic representation using abstract primitives. In *IJCAI-CANL*, pages 123–131.

[Whorf, 1956] Whorf, B. (1956). Language and logic. In Carroll, J. B., editor, *Language, Thought, and Reality: Selected Papers of Benjamin Lee Whorf*, pages 233–245. MIT Press, Cambridge, MA.

[Wilks, 1977] Wilks, Y. (1977). Knowledge structures and language boundaries. In *IJCAI*, pages 151–157.

[Wilks, 1978] Wilks, Y. (1978). Making preferences more active. *Artificial Intelligence*, 11:197–223.

[Wilks, 2007] Wilks, Y. (2007). Ontotherapy: Or, how to stop worrying about what there is. In *NLPCS*, pages 3–22.

[Winston, 1978] Winston, P. (1978). Learning by creating and justifying transfer frames. *Artificial Intelligence*, 10:147–172.