

8-1-1992

Metaphoric coherence: Distinguishing verbal metaphor from 'anomaly'

Sylvia Weber Russell

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

Follow this and additional works at: https://scholars.unh.edu/compsci_facpub

Recommended Citation

Sylvia Weber Russell - Metaphoric coherence: Distinguishing verbal metaphor from 'anomaly'. *Computational Intelligence*, 8 (3), 553-574 (1992).

This Article is brought to you for free and open access by the Computer Science at University of New Hampshire Scholars' Repository. It has been accepted for inclusion in Computer Science Scholarship by an authorized administrator of University of New Hampshire Scholars' Repository. For more information, please contact nicole.hentz@unh.edu.

**Metaphoric Coherence:
Distinguishing Verbal Metaphor from "Anomaly"**

Sylvia Weber Russell

Computer Science Department
University of New Hampshire
Kingsbury Hall
Durham, NH 03824

swr@cs.unh.edu

Published in *Computational Intelligence*, 8 (3), August 1992,
by Blackwell Publishers.

Metaphoric Coherence:

Distinguishing Verbal Metaphor from "Anomaly"

ABSTRACT

Theories and computational models of metaphor comprehension generally circumvent the question of metaphor vs. "anomaly" in favor of a treatment of metaphor vs. literal language. Making the distinction between metaphoric and "anomalous" expressions is subject to wide variation in judgment, yet humans agree that some potentially metaphoric expressions are much more comprehensible than others. In the context of a program which interprets simple isolated sentences that are potential instances of cross-modal and other verbal metaphor, I consider some possible coherence criteria which must be satisfied for an expression to be "conceivable" metaphorically. Metaphoric constraints on object nominals are represented as abstracted or extended along with the invariant structural components of the verb meaning in a metaphor. This approach distinguishes what is preserved in metaphoric extension from that which is "violated," thus referring to both "similarity" and "dissimilarity" views of metaphor. The role and potential limits of represented abstracted properties and constraints is discussed as they relate to the recognition of incoherent semantic combinations and the rejection or adjustment of metaphoric interpretations.

Key Words: metaphor, coherence, verbal metaphor, cross-modal metaphor, abstraction, extensible representations, constraints, features, analogy

INTRODUCTION

Why is the linguistic war horse “The ship plowed the sea” considered a metaphoric expression, but the sentence, “<Any object> plowed the spoon” anomalous? Or why can one “plow through one’s memories” but not “through a spoon,” even though “spoon” as a physical object may appear to satisfy more closely the selectional restrictions for the object of “plow”? The interdisciplinary literature is full of “anomalous” examples for which interpretations can (and therefore should) be found, given a particular context and a little effort. It is sometimes said that context so governs the meaning of an utterance that there is little left for the utterance itself to contribute. The “anomaly” of a linguistic expression as the term is used here, then, is a relative concept, and degrees of acceptability of an utterance are infinite. However, our initial intuitive reactions do distinguish between easily interpretable expressions and those which produce interpretations (if any) which are strained by some criterion. The determination of semantic and syntactic parsing choices, including anaphoric references, correct sense determinations, and the detection of nonmetaphoric tropes, as well as the construction of correct event sequences may depend on judgments of dependencies between concepts of such an expression. While goal planning in the sense of Perrault and Allen (1980) and Schank and Abelson (1977) can help to disambiguate potentially metaphoric expressions, the converse is also true--the probable interpretation of an expression in isolation may help to determine a plan or a goal (Russell 1986). At a more theoretical level, interpretability judgments based on the satisfaction of specific constraints could account for at least one sense in which an expression, while not rejected as anomalous, might be considered a “bad” metaphor. With the increasing accumulation of computational as well as noncomputational theories of metaphor, it seems timely to look at constraints relevant to a theory of anomaly vs. metaphor.

The above considerations suggest several questions of interest: What are the relationships which components of an expression must have with each other, in order for the expression to be characterized as a comprehensible metaphor as opposed to either anomaly or “somewhat incomprehensible” metaphor? How are extensions of meaning constrained to produce comprehensible, nonliteral language? In traditional linguistic terms, can we be systematic in indicating what types of properties are “violated” in either metaphoric extension or anomaly? How does the inadequacy of a metaphor affect its interpretation? These questions, as they apply to verb-based metaphor, can be addressed in terms of the notion of “coherence,” a concept which is interpreted variously by researchers of literal and nonliteral language, and which is perhaps most systematically described by Fass (1988). By considering how concepts combine comprehensibly, we may be able to arrive at better criteria for recognizing and comprehending metaphor in a computational environment.

Treatments of coherence in terms of semantic properties and constraints entail revisiting the linguistic question of “selectional restrictions.” The characterization of metaphor as violating certain selectional restrictions that would be satisfied by literal language is easily criticized (Ortony 1980) as inadequate insofar as it does not differentiate between metaphoric and anomalous language. Given that our concern is metaphor *understanding* rather than, e.g., why our grammatical rules are not adhered to in metaphor, we need to address this differentiation by focusing on what is preserved rather than what is violated in comprehensible metaphoric extension. In other words, both literal and metaphoric language must satisfy certain coherence criteria, i.e., selectional restrictions, in order to be termed “meaningful” rather than “anomalous.” This approach, then, is an attempt to at least partially account for metaphoric comprehensibility in terms of specific lexicon representations and their applicability to variously abstract concepts--insofar as this is computationally possible.

The domain of consideration will be confined mainly to simple sentences in which the verb is potentially used metaphorically in the immediate context of literally used nouns representing concepts dependent on the verbal concept as part of its underlying case-like structure.¹ Included in this domain are some classes

¹Semantic tension of course often spans sentence boundaries; our interest here is on potential semantic combinations of concepts, regardless of their actual lexical/syntactic arrangement. Thus we consider the format <subject verb object>, but these lexical items may actually be separated in the discourse. Similarly, consideration of comprehensible elliptical language is excluded here, though the failure to find a literal or metaphoric relation can point to its detection. See Fass (1988, 1989, 1991) for a treatment of the related phenomenon of metonymy.

of concepts normally neglected in computational practice, namely concepts referring to nonphysical entities and sometimes referred to as “abstract,” e.g., “beauty.” When a physical verbal concept is used with an object nominal in a different, nonphysical modality, we have an instance of “cross-modal metaphor.” A point of interest is the nature of similarities and differences in coherence constraints with respect to physical-domain metaphor.²

One might question forays into the area of cross-modal metaphor, when the problems of physical-domain metaphor have not yet been “solved.” There are reasons for going beyond a parochial concern with the latter kind of metaphor. From the point of view of metaphor theory, looking at a more general corpus of metaphors may keep us from missing important generalizations. The question of invariance across domains may be sharpened, for example, since cross-modal metaphor consistently filters out some types of properties which may confuse the theoretical picture for physical-domain metaphor. In addition, cross-modal metaphor motivates observations on linguistic representation of nonphysical concepts. For example, approximating a literal paraphrase of a cross-modal metaphor involves not only determining the semantics of the target expression, but also altering the syntax of the expression, since abstract nouns often conceptually represent verbal or attributive concepts.

The organization of the paper is as follows. After a brief review of some computational work as it relates to constraints on metaphor, I summarize MAP, a computational model which in its representational framework exhibits the difference between extensible and nonextensible types of components. I then consider: metaphoric coherence in terms of what happens to extensible components and the combinatorial semantics between such components and their objects in metaphoric extension; a tentative characterization of some constraints on extension together with their advantages and disadvantages; and some illustrations of how such constraints are related to metaphoric comprehensibility. A section on interpretations shows how the satisfaction of certain constraints may help to determine which metaphoric interpretation of a verb is intended. This is followed by discussion of the model as it relates to some other perspectives in relevant fields.

BACKGROUND

The question of metaphor vs. anomaly has been addressed in the various disciplines, though less so than that of metaphor vs. literal language. This research, however, has largely neglected any general semantic criteria governing the combination of *types* of concepts represented in a linguistic expression. That is, in typically focusing on the *mechanisms* of metaphor processing or recognition rather than on *representations*, metaphor researchers who have addressed verbal metaphor have generally used physical-domain examples without indicating how concepts and constraints in cross-modal metaphor might be represented. Brugman and Lakoff (1988), although (because?) they argue for a nonsymbolic approach in their research on metaphoric and nonmetaphoric polysemy, are an exception to the lack of concern for representation, though they do not frame their work in terms of constraints or go beyond topological descriptions of concepts. Hobbs *et al.* (1987), in their focus on commonsense descriptions of concepts, have also attended to the details of concept characterization (as they relate to the more important question of inferences rather than to the concepts themselves), and have formalized physical properties and relationships in terms of predicate calculus. Notably, they point out how some of these might be preserved in limited types of metaphorical usages, though they do not give any particular attention to novel metaphor or pursue the nature of constraints on metaphoric combinations in general.

Of those computation researchers who have systems which specifically address verbal metaphor processing, Indurkha (1986), Fass (1989) and Fass and Wilks (1983) include search of the source and target domains for a match (at some level) of nominals and of verb components. None, however, has devoted much consideration to cross-modal metaphor. Indurkha’s formal approach provides no rationale for his

²While in most of this discussion the assumption is that the literal domain of a verb is a physical domain, extensions of verbs from other than physical domains are also possible (Russell 1976).

constraints. Fass uses “sense frames” containing hierarchy- and structure-information to establish that, for “The ship plowed the waves,” the relevant analogy is that “ships sail through water” as “plows plow through soil,” i.e., both “move through a medium.” However, an explicit category such as “medium” (if in fact it can be located and matched in a hierarchy) is not easily applicable to nonphysical domains. We might ask what it is about a medium that *would* make it extensible. Carbonell (1982) and Carbonell and Minton (1983) present approaches for interpreting various kinds of metaphor, including some cross-modal metaphor. However, they do not attempt to judge metaphoric comprehensibility *per se*. Suwa and Motoda (1991) use a “primitive-matching” constraint to select a target-domain verb in *interpreting* cross-modal metaphor, similar in this respect to the procedure to be assumed here. However, they also are not concerned with the question of “anomaly.”

Martin (1990) bases his model on knowledge of particular conventional metaphors. Insofar as he is working from these specific metaphors rather than targeting the general nature of metaphoric extension, constraints on target roles are usually specifiable by a narrow category which closely matches or does not match; e.g., “cold” satisfies “disease” in the “give-disease” metaphor. When he addresses the task (of greater interest here) of extending his metaphors to include or learn metaphors which are similar but include a different nominal, however, he includes constraints in the form of common higher level categories of metaphors in an abstraction hierarchy. For example, his system interprets the new metaphor, “kill a computer-process” on the basis of knowledge of a known one, “kill a conversation” by referring to a “process” category which dominates both object nominals. Martin’s explicit literal categories, however, appear to be motivated not so much by coherence judgments with respect to e.g. “kill” and its objects, as by the need to establish *some* connection between these objects in seen and unseen usages. This approach is a reflection of his intentional focus on linguistic experience rather than with conceivability or novel metaphor.

The latter focus is the concern of MAP (Russell, 1976, 1985, 1986), which provides rough “literal” interpretations of isolated expressions containing a metaphorically used verb. This program (in various stages of implementation) is based on a model in which certain simple coherence criteria must be satisfied in order to interpret the expression as metaphoric (or literal) rather than as anomalous, i.e., as a candidate for other means of interpretation. Relevant parts of this program will be outlined next as a computational context for observations on coherence.

MAP

MAP as discussed here focuses on cross-modal verbal metaphor, particularly that which involves abstract concepts associated with humans, such as “joy” or “wealth.” An early pilot program processed within-physical-domain metaphor, and further research has provided a limited treatment of nominal metaphor in context (Russell 1986).

Most types of representations in AI and linguistics are not extensible, in that they are specific to the concepts contained therein, and in that there is no attempt to account for perceived analogies between concepts through similar components at some level of the representation. MAP, by contrast, attempts to distinguish extensible and nonextensible representation components of verbal and nominal concepts. According to this representation system, verbal concepts are defined in terms of abstract *extensible* structures and features and *nonextensible* content representing the conceptual domains of nominal concepts and of verbal concepts as used in their literal sense. These domains are “conceptual” in the sense that each domain is intended to represent a conceptually different modality, e.g., thinking vs. doing. The conceptual domains assigned to concepts in the lexicon, together with subcategories as used in MAP are:

MENTAL: intellect, attitude, volition

SENSORY: sight, sound, other (unimplemented) senses

CONTROL: intrinsic (talent), extrinsic (rights/duties, possession)

PHYSICAL: animate, inanimate

SPATIAL (for nominals only, e.g., “space,” “room”)

TEMPORAL (for nominals only, e.g., “time,” “year”)

The abstract concepts which constitute the extensible representations of verbs are considered to be orthogonal to or “invariant” under extension to any of the above domains. In other words, the extensible part of the meaning of the verb is that “common ground,” difficult to describe in truly abstract or “language-free” terms, which can be considered as part of the meaning of the verb in both its source domain (for which the verb usage is literal) and its target domain (the domain to which the verb is extended). The nonextensible part is the “leftover” meaning which is *specific* to either the source or the target domain. The verbal concept in its metaphoric usage therefore loses its specifically source-domain meaning and acquires some target-domain meaning, but can be understood as retaining its extensible, domain-independent structure. For example, throwing away a piece of paper and throwing away an idea can be perceived in terms of similar structures, i.e. that of leaving the state of having the paper/idea, but the actual physical and mental nature, respectively, of these events is very different. The extensible structure, then, does not reflect objective reality, but is a way of applying our cognitive or linguistic competence to talk about “elusive” concepts, particularly intangible, nonspatial concepts.

Like conceptual domains, the symbolic level of these extensible representations can also be thought of as “conceptual,” somewhat close in level to the LNR representations of Norman and Rumelhart (1975), and clearly more abstract and componential than Schank’s (1975) conceptual dependency representations, in that they factor out domain content (e.g., the P, M and A from PTRANS, MTRANS and ATRANS). Basic structures specify a (unary) attributive or (binary) relational STATE, any external AGENT that causes the state, and any TRANSITION that applies to the state (ENTER, LEAVE, or pass-THRU the state); these provide the most general and minimal level of interpretation of an expression--in the worst case MAP, in paraphrasing a metaphor, can always say at least, e.g., that someone caused the end of a mental state. Further extensible descriptors relating to these structures specify whether the state is symmetric, attempted (vs. accomplished) and potential (vs. real). Others characterize the action (if any) represented by the verb (speed, repetition, continuity), the actor (whether he/she intends the action) and any subjective response to the verbal concept (evaluation, intensity, emotional values). While the basic state structure preserves the existential content of the source and target meanings of the verb, a potential *reason* for the metaphor is represented in the verb definition by connotations expressed as a simple effect, e.g., FORCE (“effort”) for the verb construct “plow through,”³ or as a secondary structure. Other secondary structures such as a “purpose” structure are also allowed for.

Defining verbal concepts in all conceptual domains in terms of similar structures means that verb definitions reflect a metaphoric view. That is, concepts which are conceptually (literally) attributes, actions and relationships are reified and treated syntactically (metaphorically) as objects which themselves have attributes or enter into relationships. For example, “to be hopeful” or “to hope” is represented as an OBJECT (“hope”) AT a LOCATION (an animate being) in the MENTAL-ATTITUDE domain. This abstract semantic model therefore does not represent real-world actions literally in the sense attempted through systems such as Schank’s (1975). Rather, the assignment of extensible structures to concepts represents a resort to an *ABSTRACT CONCEPT AS OBJECT* metaphor (Russell 1989) in terms of which any verbal metaphor that includes a reified concept can theoretically be interpreted. In the interpretation process of a coherent sentence, the abstract structure of the verb, together with the invariant descriptors, including subjective effects, are retained, with optional translation of the verb to a verb in the target domain, i.e., the domain of the conceptual (affected) object of the expression. Any physical instruments of the process are not carried over into the target domain.

A brief summary of steps of the program to paraphrase metaphoric usages of verbs follows, using the simple sentences “The news torpedoed his hope” and “He plowed through his memories” as examples. The above listed verb descriptors, which are critical in distinguishing, e.g., “plow through” from “glide through” (see Russell, 1986) are omitted in order to focus on those components, especially the conceptual OBJECT

³In MAP, verbs are considered together with any subsequent preposition because this enables the model to recognize “X plow through Y” abstractly as a relation which can metaphorically characterize the possession of abstract concept Y by animate concept X, bringing with it connotations which might not be salient to either the verb or preposition taken separately. See Brugman and Lakoff (1988) for an alternative, systematic, purely topological approach to this type of expression.

of the verbal concept, which are subject to the constraints to be discussed later.

- (1) Identify the underlying structure of the verb/sentence by mappings from syntactic to conceptual roles, based on case-role information given in the verb definition.

The syntactic subject of the agentive verb “torpedo,” i.e. “news,” maps into the AGENT,⁴ and undergoes no further transformation in the interpretation process. The direct object “hope” maps into the conceptual OBJECT (neutral patient, theme, etc.) of the resultant LEAVE-STATE (of OBJECT BEING).

The syntactic subject of “plow through,” i.e. “he,” maps into the conceptual OBJECT, and the syntactic object maps into the LOCATION of a TRANS-THRU STATE. Since the model views abstract concepts as OBJECTS located AT animate beings, and the situation here exactly satisfies the reverse of this, the initial configuration is switched to conform to this general metaphor, so that “memories” is the OBJECT and “he” the LOCATION.

- (2) With the case-roles of the abstract verb representation filled in, we have, for the first example (in English-like representation with abstract structural information in upper-case letters), “The news CAUSED his hope to START to NOT BE.”

The second example is more complex, because of the above-mentioned reversal, the reference to different functions of the mind, e.g., Conscious Processor (CP) and Long Term Memory (LTM), and the problem of representing “going through a plural concept” as opposed to representing simply the beginning or end of a state.⁵ Briefly, however, TRANS-THRU indicates a focus on a SEQUENCE of LOCATION relationships (what happens between an ENTER- and a LEAVE-STATE), and the basic abstract structure is: “He TRANS SEQUENCE (memory AT he(CP)).”

- (3) This still metaphoric representation is mapped to the target domain: A general “de-reification” rule, <concept (noun)> AT animate-being --> animate-being <concept (verbal or adjectival form)> is used.

For the first example, the abstract structure (originally derived from the verb “torpedo”) is used to find a verb with that structure but in the subdomain of the noun “hope,” namely MENTAL-ATTITUDE, with additional primitives POSITIVE and FUTURE, giving, for “news” as the syntactic subject, “discourage”⁶ If no such verb is found (a “lexical gap” in the language--or in the lexicon), an interpretation may be pieced together by using English translations of the structural primitives, together with a “literalized” version of the reified concept. For example, the lexicon contains the mapping that the verbal and adjectival forms of the noun “hope” are (the verb) “hope” and “be hopeful” respectively. This gives “The news caused him to no longer hope.” Another paraphrase within the capability of the model, though not a literal paraphrase, retains “hope” as a noun and looks at the various conceptual domains to find a match for START to NOT BE, giving for the SENSORY-VISUAL domain, for example, “The news caused his hope to disappear” (omitted verb

⁴In a complete NLU system, “news” would actually map into part of a “causing conceptualization” (in the sense of Schank, 1975), giving a representation for, e.g., “Hearing the news caused...”

The latter two treatments have not yet been implemented with sufficient attention to generality in MAP, but the CP/LTM relationships can be assumed to be similar to those identified by Schank (1975), e.g. all mental processing is done in the CP.

⁶With such a general domain indication, other verbs of a larger vocabulary might also be matched which are less accurate. Adoption of a more specific existing or newly researched taxonomy of emotions/attitudes is planned.

descriptors give “suddenly, with intensity”).

For the second example, “memory” indicates that the target domain is MENTAL-INTELLECTUAL. More specifically, “memory” is an OBJECT in the LTM “part” of this domain, which in being acted upon becomes successively located AT the CP part of “he.” A target candidate for a CP verb is “consider,” if the word “memories” is retained in the output paraphrase. The paraphrase is then “He considered his memories” (omitted verb descriptors give “slowly, continuously, with effort”). Another paraphrase alternative includes “remember <content of what is remembered>” as a verb which explicitly represents the location in the CP of an “object” previously in the LTM.

As MAP is oriented toward breadth rather than target-domain detail, the paraphrases produced often seem too general. They appear, however, to be close in content to the types of responses Gentner and France (1988) observed empirically, which they classify as “minimal subtraction” (of meaning from the verbal concept in its original sense).

METAPHORIC COHERENCE

The task before us is to find a means of distinguishing coherently metaphoric usages of verbs from usages which are neither literally nor metaphorically coherent. Informally, this is a matter of determining which verbal concepts can “go with” which objects. One traditional means which suggests itself is to define a set of properties, features or categories to describe objects and the constraints imposed by verbal concepts on those objects. In most feature-based theories of metaphor, there is a set of descriptive features, some of which may be “violated” in metaphor and some not. Theories of metaphoric *incoherence* are even more vague, though to some extent perhaps necessarily. Thomas (1969), for instance, says “When *too many* features are incompatible, a metaphor simply becomes preposterous” (italics added). Most AI models of metaphor which search and match concepts in different domains can interpret coherent metaphors, but leave the determination of incoherent ones to default. In an attempt to provide more specific criteria for judging such expressions, two types of properties in MAP, “conceptual domains” and “conceptual constraints,” determine the recognition of a *possible* cross-modal metaphor vs. a literal expression, and the recognition of a coherent metaphor vs. a nonliteral, nonmetaphoric expression respectively.

Domain Consistency and Metaphoric Extension

A necessary though not sufficient principle governing recognition of cross-modal metaphor is a lack of “conceptual-domain consistency” in a literal reading of the sentence.⁷ The assumption here is that a verbal concept and its possible associated objects are semantically interdependent—they share each other’s properties. A verb used in one conceptual domain with an object in another is thus literally incoherent; physical concepts do not literally mix with mental ones in an utterance.

This inconsistency is resolved when the verbal concept is comprehended as extended to the domain of the nominal.⁸ In linguistic terms, we could speak of domain consistency as a fundamental “selectional restriction” that is a minimum necessary for coherence. For a simple expression in which one or more nouns provide the immediate context for a potentially metaphoric verb, this principle can be a means of recognizing nonliteral language; if the verb in its *literal* sense and the conceptual object of the sentence are inconsistent in domain, as in “siphon off the idea” or “his pride broke,” the expression is not a literal one.⁹

⁷This approach does not imply that metaphor is to be thought of as “deviant” with respect to literal language, or sought only after literal interpretations fail, but rather as a different kind of predication with its own set of constraints (Russell 1986).

⁸It could of course be the case that the nominal concept is extended to the domain of the verbal (see Levin, 1977, for comparisons of these), but that is not of interest here.

⁹Conditions of course also apply to the actor of the verb. Broadly, an actor in a MENTAL, SENSORY, CONTROL or PHYSICAL-animate domain must be animate for the expression to be literal. Otherwise we potentially have a personification, as in “The trees began to think about dropping their leaves” or as in *one* interpretation of “My car drinks gasoline,” an example pursued at length in earlier work by Fass and Wilks (1983).

If, further, the verb is seen as extended to the domain of its object nominal, this extension is “conceptually coherent” providing the verb and its dependent nominals combine in a “conceivable” way (though not necessarily a “usual” way), according to any selectional constraints. If the verb is thus “comprehensibly” extended, the expression is a metaphoric one in that domain of the nominal; otherwise it is anomalous, i.e., the object cannot serve as the “real” conceptual object of the verb. A comprehensibly extended verb can be seen as analogous to the verb in its literal sense, the (abstract) structure being the same but the domain different.

Literal domain inconsistency as a possible indicator of metaphoric extension of a given verb as modeled by MAP invites the question of what happens to the meaning of verbs and nouns in metaphoric extension, and therefore to their combinatorial semantics. The empirical research conducted by Gentner and France (1988) addresses this question. They discuss subjects’ interpretations of semantically strained sentences in isolation in terms of “verb mutability,” i.e., the potential of the verb to change in meaning to adapt to the domain of the object. They note that mutability may account for polysemy (at least some of which, it would seem, is due to frozen metaphor). This appears reasonable. In terms of models such as MAP, however, as well as in terms of Gentner’s and France’s observations that abstract relations are extended to interpretations, whereas specific objects are not, there may be a co-existing reason for both verb mutability and polysemy. It could be that a verb can be understood in so many different contexts, though changed in meaning, precisely because a core component of its basic meaning is *immutable* and is therefore recognized in whatever conceptual domain it is used. The fact that the verb submits to the noun in adapting its meaning thus in no way contradicts the centrality of the verb as the authors suppose; structurally, it is still the verb which determines the expression of the thought to be conveyed. Rather than downplaying the verb, Gentner’s and France’s studies, as well as elements of MAP and other programs, appear simply to reflect both *dissimilarity* and *similarity* views of metaphor, namely that the domain changes but some part of the structure must remain the same. It is the “similarity” aspect which determines the extension of constraints as well as of structure.

Conceptual Constraints

Domain consistency through extension does not in itself guarantee a metaphorically coherent expression. Some kind of constraints are needed to relate constituents of the expression in a way which satisfies the analogy perceived or imposed by the originator of the expression. Constraints may be expressed as features describing static characteristics of a constrained object, a structure describing the use or function of an object, or a category or categories in which these features and/or functions are implicit. One application of constraints is illustrated by Fass (1989), who, as mentioned earlier, shows how “The ship plows (or plows through) the waves” can be interpreted metaphorically through a sense frame search which determines that both a ship and a plow “move through a medium.” However, explicit categories on the level of “medium” do not help in comprehending extensions to nonphysical domains, such as “The cashier plowed (or ‘plowed through’) her memories.” The difficulty of application is at least partly due to the fact that such categories include domain-specific content, which is not extensible.

While Martin’s (1990) model differs from that of Fass, it appears that similar observations on the nature of constraints would apply. Martin’s subcategorization of “process” in his *NONLIVING THING AS LIVING THING* hierarchy seems reasonable, but does not reveal what in general constrains that metaphor. To “kill the lights,” for example, is a common conventional metaphor, yet “lights” apparently does not fit into Martin’s hierarchy (an overconstraint) except possibly at the top level (*NONLIVING THING AS LIVING THING*), which is no constraint at all (and which would be irrelevant, e.g., to the similar metaphoric expression, “torpedo the process”). The point here is not to say that there should be a category to include “lights,” but rather to suggest that we may want to ask what it is that is salient about concepts with respect to killability, i.e., what makes them seem life-like. That is, what type of concept could conceivably fit into any given one of Martin’s metaphor-map hierarchies? In psychological terms, how might humans understand metaphors based on conceptual and semantic knowledge of a concept as opposed to knowledge of previous metaphoric usages of that concept?

With this question in mind, the concern in this discussion is a model which might be enlightening as to a more basic metaphoric competence which relates the understanding of *any* simple-sentence metaphor

to a perception of the constituent concepts. Further questions motivated by this approach might be: Are metaphoric objects of “kill” those which appear to have life because they have an effect on us, or perhaps because they have “meaning”? What does this say about subjectivity in metaphor? These are theoretical questions, but are not without practical relevance to future computational models. It would be interesting to know *why* conventional as well as novel metaphors work. Characterizations of constituent concepts and constraints will not in themselves answer the above questions, but they provide certain properties which can be further analyzed and discussed.

The focus in this discussion is mainly on constraints between verbal concepts and their objects. Constraints by verbs on dependent object nominals depend directly on metaphoric extension of components of the verbal concept. In terms of the previous section, some constraints on objects appear to be extended along with immutable components of the verbal concept, i.e., components which are perceived as similar in the source and target usages of the verb. For instance, there is only so much we can say about the verb “plow through” that is immutable or invariant, without speaking of its conceptual object. We can speak of transition, force, and continuity as immutable components applying solely to the verb in both literal and metaphoric senses. However, the involvement of an object (if it exists) of a verb is also part of the verb meaning, and the constraints which say what can coherently be the object in an *extension* of the verb are therefore also invariant under extension. If a literal constraint imposed by “plow through” on its object is that the object have a certain texture or “composite” character, then some extensible part of this constraint must also hold for the extended usage of “plow through,” and, as for components of the verb itself, some domain-specific meaning of the constraints must be left behind.

Features. We discuss constraints here mainly in terms of conceptual attributes or features, though it is not being claimed that features are the most important part of the semantic description of a nominal (Russell 1986).¹⁰ The purpose of extensible features of this kind, based to some extent on “abstract naive physics/mathematics,” is to determine semantic relations which are *conceivable* though not necessarily usual or expected. In the physical domain, descriptors at this level allow interpretation of examples such as “plowed through the trash” (even though “trash” would not be expected to be defined as a “medium”), because “trash,” in being a collection of parts as opposed to a solid unit, is +COMPLEX and +FLUID.

For cross-modal metaphor, such features, since they apply to concepts which are “objects” only in a metaphoric sense, are themselves metaphoric. For example, “plowing” an abstract concept implies thinking of that concept as metaphorically +COMPLEX, i.e., not elementary or simple. As a constraint, this feature value would be satisfied by a noun representing a composite event, e.g., a “concert,” but not a simple action, e.g., a “throw” or attribute, e.g., “truth.” This means that in specifying abstract features which will “work” for cross-modal metaphor, knowledge base editors have to do some metaphoric thinking, and perhaps expect a lack of human consensus on what a particular feature might mean. As Tourangeau and Sternberg (1982) note, features cannot be identical for concepts in different domains; at best they are analogous.

Some relevant *literal* nominal features, which emerged together with early work (Russell 1975) on a taxonomy of naively perceived relationships for purposes of phrase disambiguation are:

¹⁰The term “features” is used loosely to mean some formalized property which may apply in combination with other properties. While some features in MAP simply have “+/-variable” values, the CONTAIN property associates with a “+” value a further simple specification as to what *type* of concepts can conceivably be contained by the nominal. Others are or could be 3-state variables. Of greater importance is that the set of semantic descriptors be kept small and not open-ended.

+/- SHAPE (vs. amorphous, mass)
 1-DIMENSIONAL (linear-like)
 2-DIMENSIONAL (area, surface-like)
 PART (of something)
 FIXED (attached to something)
 CONTAIN (2 types: surround or embed)
 COMPLEX (vs. elementary)
 FLUID (liquid, gel or particulate vs. solid)
 ANIMATE

In addition, the FUNCTION of a nominal, which can both serve as a constraint and enter into interpretations (see section on constraint determined interpretations below), is both expressed as a feature and expanded in terms of an abstract verbal representation.

The level of the conceptual characteristics observed and involved in developing this set is similar to that of the “cognitive notions” which Talmy (1978, 1987) discovered by observing grammatical specifications. For instance, MAP’s SHAPE feature corresponds to Talmy’s perhaps more aptly named “boundedness” feature. His feature set as a whole differs somewhat from MAP’s, partly because he is concerned not so much with semantic dependencies and constraints as with generalization of his cognitive notions over both nominals and verbs, with emphasis on grammatically determined structural notions.

Talmy’s treatment of both nouns and verbs as subject to the same cognitive notions (or feature set) makes his approach particularly interesting with respect to our focus on extensible features. One could view his approach as treating features of actions (lexicalized as verbs or verb constructs) as metaphoric extensions of features of matter (lexicalized as nouns), given that the appropriate grammatical elements are present. For example, as an illustration of the “plexity” category, he designates as “uniplex”:

MATTER: A bird flew in.

ACTION: He sighed (once).

He designates as “multiplex”:

MATTER: Birds flew in.

ACTION: He kept sighing.

Here the grammatical elements “kept” and “-ing” render the verbal concept “multiplex” similar to the way the element “-s” makes the nominal concept “multiplex.” The action “kept sighing” is not literally plural, but we apparently cognize it as having a grammatical component in common with the concept of “plural.” The feature “multiplex” which represents this component therefore applies analogously and abstractly to the verbal as well as the nominal form. Talmy’s identification of categories/features which appear to apply similarly to different conceptual or grammatical categories is of relevance to a determination of a more definitive feature set for the somewhat different purpose of MAP, and suggests further research in this area.

Given that we have such a physical-domain feature set, how do the members of a feature set extensible to abstract concepts compare with those of the physical-domain set from which it is drawn? Neither set in MAP is well-developed. However, we can say about the relationship between the two types of sets that, just as physical aspects disappear from a concept extended to a nonphysical domain, physical features merge (are mutually redundant) in nonphysical domains, because of the irrelevance of distinct topological characteristics. The effective abstract feature set can therefore be expected to be smaller than a literal feature set. For purposes of discussion, a sample extended set as derived from the above literal set consists of a subset interpreted abstractly:¹¹

¹¹Some of the problems of defining conceptual features are similar to those involved in defining conceptual attributes of objects in abstract verb representations. No definitive set of primitive attributes has been established for MAP. It can be said, however, that incidental attributes such as “color” and particular shapes such as “pointed” are not considered primitive, but “intensity” and “evaluation” are (see Osgood, 1980, and Aarts and Calbert, 1979, for evidence on the metaphoric extension of such attributes).

+/- SHAPE (vs. amorphous, mass)
1-DIMENSIONAL (linear-like)
FIXED/PART/CONTAINED (subordinate)
COMPLEX (vs. elementary)
FLUID
ANIMATE (dynamic)

Observations about Features. Several observations can be made about abstract features based on a set such as this:

- (1) Physical details become irrelevant;
- (2) Distinctions between some features are lost;
- (3) They may underconstrain interpretations by their general nature;
- (4) They may overconstrain them in nontypical cases;
- (5) They may overconstrain interpretations because of human misjudgment and variation in defining them as constraints.

(1) Physical details: My focus differs somewhat from the approach of Hobbs *et al.* (1987), who in their formalizations of physical properties and their implications note that the property of being wedge-shaped concentrates a force. This level of detail is omitted from our set insofar as particular physical shapes do not appear to be easily carried over to nonphysical domains in prosaic discourse, although given a context which explicitly paraphrases or elaborates on the metaphor, any prototypical physical detail could conceivably become relevant.

(2) Lost distinctions: In the *extended* feature set, literal features which differ only in some component physical aspect may merge in nonphysical domains. Thus PART (of), FIXED/ATTACHED (to) and (be) CONTAIN(ed) (in) merge into a vague “subordinate” concept labeled by only one of these feature names. A 2-DIMENSIONAL, flat or area-like quality does not mean much in nonphysical domains and thus also is irrelevant as a feature.¹²

(3) Inherent underconstraint: Both the basic and the extended sets, being of small size, are liberal in allowing computational interpretation of novel metaphor prevalent in various kinds of discourse. A corresponding disadvantage of course is that such features may underconstrain interpretations. This disadvantage is somewhat offset by the earlier observation that conceivability judgments carry only a part of the burden of interpretation.

(4) Inherent overconstraint: This may result, when only typical views are adhered to in the determination of constraints. Most people can think of examples which are “comprehensible” in spite of apparent violation of constraints. This problem is discussed in the section below on constraints and comprehensibility.

(5) Faulty constraint specification: Besides the problem of judging the relevance of extended literal features, there is a potential problem of using as a constraint for a PHYSICAL-domain usage a feature value which is *usually* applicable in physical-domain usages, but is not *necessarily* imposed by the verbal concept. This may result in a false constraint by the verb in an extension to another domain. For example, an early version of MAP specified +FIXED for the object of “plow.” While the usual literal object of plowing will certainly be +FIXED, this constraint is not dependent on the nature of “plowing” and thus should not be specified as an extensible constraint. If it is, then examples such as “plowing memories” or other

¹²Note, however, that a stereotyped subjective *effect* of being “flat,” namely a negative interest, is carried over into nonphysical domains by components of the abstract verb representations as described in the section on MAP.

conceptual objects will be wrongly hypothesized as anomalous. It seems that extensible components and their use are particularly susceptible, not only to idiosyncratic understandings of word meanings, but also to undesired consequences resulting from inadequately considered representations.

Constraints and Comprehensibility

Given the necessarily vague nature of judgments on the applicability of such features to a concept represented by a noun, are they of any value in distinguishing metaphor from anomaly? As an example, the verb construct “plow through” imposes the constraint +COMPLEX on the object of the plowing. A “plural” concept will satisfy this. This constraint would appear to account for the observation that “She plowed through her memories” is less strained than “She plowed through the (musical) note.” Similarly, even though “plowing through memories” extends a physical action to a mental domain, it is less strained and more conceivable in isolation from discursive context than “plowing through a (-COMPLEX) spoon,” because the invariant part of the meaning of COMPLEX applies to “memories,” but not (in the current version of MAP) to “spoon.” While “plow” has an interpretation somewhat different from that of “plow through,” constraints on the object are similar; therefore “plowed the waves” and “plowed her memories” (changed their “shape”) are both more comprehensible than “plowed the spoon.”

As another type of example, for “The farmer plowed through the book,” conceptual features could be used to select the MENTAL aspect of “book” rather than the PHYSICAL aspect, despite the strong association of the literal usage of “plow” with “farmer.”¹³ Disambiguation at this intra-sentential level is due in part to the fact that the MENTAL aspect of “book” is defined with a +COMPLEX feature value, satisfying the conceptual selectional restrictions of “plow through,” whereas the PHYSICAL aspect of “book” is -COMPLEX, and in any case does not satisfy another feature constraint on “plow through,” namely +FLUID.

In both of the preceding examples, seeing an object such as a spoon or the physical aspect of a book as -COMPLEX reflects an intuitive, naive characterization of the object, since in a strict analytic sense, almost any object is +COMPLEX. Consider the example

The electrons plowed through the spoon¹⁴

Here the actor “electrons” determines a kind of plowing which is much smaller in scale or finer in granularity than that thought of in connection with visible actors of the everyday world. A -COMPLEX constraint on the object of “plow through” would block this metaphorically coherent sentence. There are at first consideration at least three ways of approaching the deviation suggested here:

- (1) The example enters a scientific world not subject to the naive perceptions which guide the constraints on physical concepts in everyday language. This is not to say this statement should not be interpreted, but that it is outside of the incomplete but large framework used here.
- (2) The determined constraints depend on one perspective. The system should be expanded to allow for shifts in magnitudinal perspective, like “zooming” in or out. This would be similar in concept to Talmy’s (1978) identified cognitive process of “magnification” or “taking the close-up view.”
- (3) Constraints on objects always are subject to the immediate context provided by the actor of the verb. By “context” is meant here not some vague component of the meaning given by preceding discourse, but rather the relationships between the object and other nominals *through* the verb. That is, ships can plow a sea only if ships can have some locative relationships with a sea in which the ship is at the same or smaller order of magnitude than the sea. In the above example, we have the unusual situation that the “plowing” concept is at least an order of magnitude smaller than an

¹³If a nominal belongs to more than one conceptual domain, each domain specification has its own “track” of features. An expression including such nominals, such as the present example, could also be viewed as a “frozen metonymy” (see Fass, 1983, for alternative resolution).

¹⁴I would like to thank an anonymous reviewer for this example.

already “small” object, i.e. a spoon. This comparison or relation has not been implemented in MAP, but should be considered in refining it.¹⁵

While such considerations are of semantic interest, we might ask whether constraints on nominal objects of verbal concepts play a significant role in the AI goal of text comprehension. Cases in which humans could not resolve a metaphorically ambiguous sentence through extra-sentential context are probably infrequent. In practice, however, successful computational use of such context is not easy, and the above constraints do serve as an aid. One of the more practical applications of such constraints could be in the determination of antecedents of an anaphoric reference. Consider:

The rifle was found in the trash left on the ground around the pole.
John plowed through **it** and reported his discovery to the police.

There are various ways of determining antecedents of pronouns in sentences of this type, using both semantic and world knowledge (see for example Wilks, 1975; Hobbs, 1977). However, a system with no knowledge of metaphoric constraints would have difficulty in selecting *any* of the possible nominals of the sentence as the object of “plow through.” Here we can use this kind of knowledge in terms of the above features to reject “rifle,” “ground” and “pole” (all -FLUID and the latter two -COMPLEX) as the object of the metaphoric “plow through” and accept “trash” as the antecedent of “it.” A similar example with abstract nominals is:

The discovery of the rifle provoked a great deal of discussion.
Plowing through **it** was an arduous task.

Here “discovery,” a -COMPLEX action, and “rifle” are rejected as the object in favor of “discussion,” which is +COMPLEX and +FLUID.

It can be seen that the “constraint relaxation” (if it is seen that way) underlying metaphor vs. literal usages does not depend on an undifferentiated feature set in which +COMPLEX, for example, is given the same status as +MENTAL. While no claim is being made for the psychological reality of the individual components presented here, it is evident that the properties referred to can be classified in *some* way which shows analogical relationships between constraints as well as between the verbal concepts themselves. The above hypothesized verb-object combinations suggest that while literal coherence depends on domain- as well as structure-based properties, metaphoric coherence is independent of domain specifications *per se*. “Plowed his memories” is metaphorically coherent, and “plowed the note” is metaphorically incoherent, even though there is a domain “violation” with respect to “plow” in both cases. In a parallel vein, “the ship plowed the sea” is metaphorically coherent, and “the farmer plowed the spoon” is metaphorically incoherent, even though “sea” and “spoon,” being both physical concepts, may seem closer in meaning than “memories” and “sea.”

CONSTRAINT DETERMINED INTERPRETATIONS

An initial failure to satisfy constraints for a coherent metaphor in the process of text interpretation does not necessarily imply that some other parse or a different kind of trope should be considered. Another possibility is that an alternative but still metaphoric interpretation should be sought. In other words, verbs may sometimes be interpreted in different, though perhaps related, metaphoric ways, depending on what types of selectional constraints are satisfied by the object nominal. Examples for comparison are

(1) “His efforts sharpened his mind”

¹⁵The problem is analogous to one considered in Russell (1975), namely in order to know whether “I saw the birds flying to the convention” can be interpreted as the birds flying to the convention, we have to know whether it is conceivable or metaphorically reasonable for birds to be “at” a convention. In other words, a perceived problem of context can be reduced to the question of a simple conceptual relation or dependency between concepts.

(2) “The event sharpened his memories”

Here we have (at least) two different cases of what might be salient, corresponding to whether we think of the imagined object of sharpening as a cutting instrument which is potentially used to operate on objects (1) or as a general object simply to be perceived (2). Minimal, relatively literal interpretations of the above examples that we might want to target here are:

- (1) “His efforts caused his mind (i.e., him) to function (i.e., think) better”
- (2) “The event caused him to perceive or distinguish his memories better”

Informal literal definitions of “sharpen” upon which the above interpretations could be based read something like:

- (1) “cause something to be sharp for the purpose of cutting well”
- (2) “cause one to perceive something better”

If we view the transitive verb “sharpen” as having these literal definitions, we can use corresponding abstract representations in combination with constraints to arrive at the above interpretations. A general English-like paraphrase of an abstract extensible representation comprising these two aspects of “sharpen” is:

person_i (or event or situation, i.e., the “sharpener”)

causes object_i (knife, mind, memories) to be more <attribute_i value> with

- (1) the purpose that
person_j use object_i (+FUNCTION) to <FUNCTION of object_i>
better
- (2) the effect that
person_j perceive object_i
better

In both cases, person_j may = person_i.

Interpretation (1) says that, in order for it to apply to a usage of “sharpen,” the object must have an inherent FUNCTION.¹⁶ This is satisfied by “mind,” (giving a target interpretation in terms of thinking better), but not by “memories.” For interpretation (2), there are no constraints on the object, so either “mind” or “memories” *could* be interpreted this way (i.e., in terms of being more clearly distinguished), but the interpretation corresponding to the satisfaction of the more specific constraint is preferred.

Constraints, then, determine whether the direct object in the above sentences maps into an “instrument” which functions better as a result of the sharpening, or into a “passive object” which is better distinguished. Different interpretations of the verb, for metaphoric as well as literal language, thus derive from determination of which kind of object we are dealing with. Just as physical sharpening can refer to a functional use or simply to a physical property, nonphysical or abstract sharpening can do the same. In both types of domain, the constraint +FUNCTION must be satisfied for the functional case.

In the attempt to resolve an ambiguity of this nature, it becomes evident that there may be limits to our ability to represent all the subtle ways in which an abstract concept, particularly, can be metaphorically structured and interact with other concepts. Whether to “sharpen their rights” means to make existing rights more intense or operative on one hand, or to make clearer the *specification* of those rights (an embedded metonymy) on the other, may not seem to be problematic for a human, but the somewhat different interpretations resulting from a precise analysis of this phrase may entail different inferences.

What if no set of constraints is satisfied, but an interpretation is required because of other text processing considerations? One approach is to provide an interpretation in which only the most basic, i.e. the existential parts of the verb structure, enter into the interpretation, under the assumption that this is simply

¹⁶While the literal PHYSICAL interpretation of the transitive verb “sharpen” requires its object to have a point or an edge, no feature corresponding to this is obviously extensible to nonPHYSICAL domains. This is an example of nonextension of detail as described in the previous section on coherence.

an unrecognized idiom or a “bad” metaphor. If the only available parse is one which represents “the farmer (or even the ship) plowed through the spoon” (corresponding to the situation in which humans would find an interpretation if forced), a possible, minimal interpretation is that the “plower” was in motion, made contact with the spoon and ended up somewhere else. “Plowed the spoon,” also initially judged as anomalous, would receive a default interpretation saying that something was done to the spoon to change its character in some way.

RELATION TO OTHER PERSPECTIVES

In terms of traditional linguistic theory, the presented way of judging metaphoric coherence maintains the distinction between violated and *preserved* types of features, a characteristic missing from feature-based systems such as that of Levin (1977), which in other respects has similarities to some aspects of MAP. For instance, Levin shows how the example “His courage evaporated” might be construed in terms of the transfer/addition of an “Abstract” feature into the selectional restrictions of “evaporate.” Levin treats this transfer as automatic, with no judgment as to how comprehensible this nonliteral usage is. He simply adds a feature to the allowable list when a “deviant” usage is encountered. This is done even when the expression conceptually does not make sense, as in “hopeful sign,” an expression which appears to have originally represented a speaker-convenience-related transfer of “hopeful” from the “hoper” to the object which instilled the hope, rather than an extension of those concepts which can hope. For Levin the verb in a verbal metaphor is not *subject* to the constraints expressed as features of the nominal; rather it *assumes* these features.

In MAP, the discrepancy in “Abstractness” between literal usages of “courage” and “evaporate” would simply be indicative of an extension to another domain, the comprehensibility of which is determined by the satisfaction of extensible constraints. This system is more explanatory than approaches which simply note that a feature has been violated, in that it takes into account the understanding and generation of analogous usages through abstraction. That is, the fact that some part of the structure is preserved under extension reflects our basic ability to abstract and thus to identify how certain verbs (cf. Jackendoff 1976) appear to be analogous. The orthogonal nature of the relationship between nonextensible domain-related features and extensible structure-related features helps to explain why, as asked earlier, the more “radical” extension in cross-modal metaphor may be more comprehensible than a within-domain metaphor.

How features are used in the two systems is closely tied to the presence or absence of a “conceivability” criterion. The notion of conceivability I assume is somewhat different from that of Levin, though both approaches make the distinction that the absurdness of a situation does not *necessarily* imply inconceivability. Levin’s neglect of conditions of transfer/extension is related to his idea of metaphoric conceivability. Since he does not distinguish between various kinds of poetic and prosaic contexts in his construals, anything is conceivable, so features are always transferred. However, in allowing that *anything* is conceivable, Levin has no means of distinguishing between metaphor and other occurrences of nonliteral language. I would maintain that when the purpose is not poetry but ordinary discourse (which even with novel usages appeals to the reader’s model of the world for coherence), immediate dependencies between e.g. physical and nonphysical concepts are not conceivable unless certain constraints deriving from the abstracted meaning of the verb are satisfied. If such constraints are not satisfied, coherence should be sought elsewhere.

One could argue that conceivability is inherently fuzzy and that drawing the line is a futile endeavor. As Smith, Rips and Shoben (1974) and others have observed, category membership (and one could add, feature attribution and constraint) is often a matter of degree. Drange (1966), however, in considering *literal* conceivability, maintains that the “thinkability” or “conceptual meaningfulness” of a proposition is fuzzy only insofar as there is no clear consensus about what the constituent concepts themselves mean. Vagueness here may come from the incidental composition of a concept (e.g., whether an unusual extension of a bed may have a bench attached, making “the seat of the bed” thinkable), or from a verb’s crossing the threshold from a literal to a metaphoric usage (e.g., that “squirrels, or even bacteria, like coffee” is literally thinkable, but that “plants like coffee” is only metaphorically thinkable. Drange thus shifts the burden of drawing the line of conceivability from the coherence of a proposition to the individual concepts.

While Drange's claim about drawing the line of literal conceivability is subject to argument, the existence of a *metaphoric* conceivability dichotomy is necessarily assumed here, since interpretation choices must be made. Of course, this does not mean that we can arrive at a satisfactory means of "line drawing." In considering metaphoric as opposed to literal conceivability, we have to admit not only that the constituent concepts (particularly abstract concepts) might be vague or ill-defined, but that while the literal unthinkability of a proposition itself *may* be clear, the thinkability of the proposition in a metaphoric sense may not be. Furthermore, various extensions of a verb, together with other elements of the discursive context, may interact with dependent nominals and affect *their* interpretation. It is no wonder that Levin is willing to simply allow everything. However, in building an actual system which distinguishes metaphoric and literal language from other modes of interpretation, we must link a conceivability limit to some criteria. The abstract features we employ to express these criteria may themselves be subject to variation in interpretation, but once specified as properties or constraints, they commit MAP to a decision, either that the expression is metaphorically coherent as it is, or not. In the latter case, other facilities of the NLU system can attempt an adjustment, such as the detection and interpretation of an embedded metonymy or ellipsis, as in Drange's "seat of the bed" example, or accept that the expression is a misparse.

The MAP constraint system used to establish the conceivability in question both builds on and differs from Drange's theory of metaphor. Drange refers to metaphor as a type crossing (though of course not the converse). This corresponds to MAP's assumption of domain inconsistency as an indicator of nonliteral language and as a necessary condition of metaphor. Like Levin's features, however, Drange's types conflate what we have called domains and structures (and features). Maintaining the distinction between domains and structures can help to explain literal and metaphoric coherence in a way that an undifferentiated type definition cannot: A type crossing in which object nominals do not satisfy constraints imposed by the extensible verb structure is neither literally nor metaphorically meaningful, regardless of domain; a type crossing in which the nominal does conform, but the domain is different is metaphorically meaningful (in the sense of thinkable). An expression in which the structure conforms and the domain is the same, i.e., in which there is no type crossing, is literally meaningful.

MAP's concern with conceptual relationships between concepts as a constraint on both literal and metaphoric coherence indicates the relevance of "naive physics" models. As in the domain of literal language, naively perceived physics- or mathematics-like structural descriptions of concepts and their relationships play a large role in how one concept can be associated with another in various kinds of knowledge representations. While the discussion here has been restricted mainly to dependencies between verbal concepts and their objects, research on such models is of interest because research on metaphoric coherence must build on them to arrive at *abstracted* versions of these relationships in terms of constraints and features. With the addition of the metaphoric dimension, a model of these relationships should, as in the area of literal language understanding, hold for "all" subworlds. As Hayes (1985) argues, it is of little value to restrict one's research to apparent success within a toy domain; systems should aim for breadth, which in turn will make evident potential areas of development in the sense of density.

MAP at this stage makes no pretense about density. However, its concern with breadth as well as uniformity of application is illustrated by its venture into nonphysical conceptual domains. MAP not only models relationships which hold within all of these subworlds (e.g., "learning" is a TRANSITION of a MENTAL-INTELLECTUAL OBJECT which results in a MENTAL STATE, e.g. one's having the OBJECT in mind), but, more notably, relates concepts in one subworld to concepts in another through extensible structures ("learning" has an abstract structure similar to that of "acquiring" or "inheriting," and "having something in one's mind" has the same structure as "having a physical possession under one's control." Looking at the world in terms of "all" conceptual domains is illuminating for cross-modal metaphor, in that it in fact can lead, as Hayes suggests, to potential areas of development. One area might be a treatment (in an other than *ad hoc* fashion) of what relationships might *not* hold across domains. For example, ideas can be "sent," but do not leave their source as physical objects do; there *are* domain-dependent discrepancies in extensible properties and inferences generally hypothesized in MAP. The current version of MAP would (perhaps justifiably) have trouble with a sentence such as "The idea she tossed out (-FIXED constraint) to me remained rooted (+FIXED constraint) in her mind." Narrower looks at individual domains as well as the meanings of features in *each* domain can help to fill in the gaps. Barnden (1990, 1991), for example,

explores reasoning with metaphoric structurings of mental states and their inferences. A model such as MAP, with its broader view, can for its part provide a guide for considering similar metaphoric structurings for other types of concepts.

With its concern of distinguishing novel metaphor from incoherent expressions rather than recognizing known metaphors, my approach differs from that of Lakoff (1986) and Lakoff and Johnson (1980) in the way it handles different metaphoric structurings. If MAP is given the expressions “He boarded up his mind” and “He sharpened his mind,” it works only from its knowledge of the constituent concepts in arriving at a metaphoric understanding of the expressions. “Boarded” requires the object to be able to CONTAIN, for the interpretation that he did not let any new concepts into his mind, i.e., start thinking about something new, and “sharpened” requires the object to have a FUNCTION as above. This procedure is equivalent to Lakoff’s and Johnson’s (presumable) characterization of these sentences as linguistic instances of *MIND AS CONTAINER* and *MIND AS TOOL* metaphors respectively. MAP, however, uses general semantic and conceptual knowledge, rather than relying on previously encountered metaphors.

This approach can be seen as complementary to the computational approach of Martin, which uses knowledge of conventional metaphors like those of Lakoff and Johnson. Martin’s approach works “upward” toward generality from lower-level metaphors which, since they have previously been encountered and stored, support the likelihood that a metaphor recognized by his system is in fact intended. That approach is a useful way to start dealing with mundanely metaphoric text, if generalizations are uncovered along the way. The approach presented here *starts* with a more general metaphor, the *ABSTRACT-CONCEPT-AS OBJECT* metaphor, which generalizes Martin’s (1991) “core mappings” (e.g., “beliefs are objects with locations”) as “all attributes/acts/states of animate beings are objects with locations” and lets metaphoric usages such as those based on “beliefs are objects” be instantiated. MAP thus provides a model based on conceptual/linguistic principles, in terms of which Martin’s observed metaphors can be understood and theoretically predicted.

The linguistic assumption here is that the *ABSTRACT CONCEPT AS OBJECT* metaphor underlies reification--the syntactic device through which we speak of conceptual (in the sense of Schank, 1975) attributes, acts or states as if they were objects, i.e. as nouns. These “objects,” whether their semantic domain is that of e.g. mind or power, and whether conceptually they are e.g. attributes or acts, metaphorically enter into the same case-like or logic-like relationships as physical objects do. The metaphoric viewing is therefore at the linguistic/conceptual level rather than at the lower, semantic-category or -property level. The process of selecting a target interpretation involves undoing the *ABSTRACT CONCEPT AS OBJECT* metaphor through selection of a dereified target concept.

The psychological assumption behind working through similar abstract structures rather than known metaphors could be described as follows, given e.g. “The news torpedoed his hope”: An originator of the metaphor, who suddenly loses (or is writing about someone losing) hope is reminded, through both emotional and perceived structural similarity, of his hope (or, metonymically, of himself) being torpedoed. He may or may not also know a literal way of describing this, but does have the sense that a “destruction” inference (of “torpedo”) applies here. In another person the same experience might lead to a metaphoric usage of a different verb also expressing destruction, which MAP could also interpret, given the abstract definition of that verb in the lexicon. It thus discovers metaphoric usages that Martin’s system would already have to know about. That is, Martin would need a “hope as boat” metaphor, a type of knowledge which has little to do with metaphoric competence, or perhaps a metaphor at a somewhat higher level in a hierarchy. To have a flexible system, he would need to enter in an explicit metaphor for every reminding or type of reminding that a person could have, related to the topic she wishes to describe.

Martin’s focus on specific instances or themes of a metaphor could be advantageous in providing elaborations on a metaphor, and in any case is no doubt more robust than MAP in its detail. MAP, however, is an attempt to provide a more general account of metaphor understanding and generation in terms of the ways we might think metaphorically about actions, attributes and states, and syntactically express predications about these entities. This account should theoretically be able to understand conventional as well as novel metaphors of the type discussed here, given that appropriate semantic descriptions of these concepts are present in the lexicon. In a sense the relation of MAP to Martin’s system is analogous to that of the use

of plans and scripts in discourse understanding, in that “plans are where scripts come from” (Schank and Abelson 1977), and in that scripts may be more reliable, *if* they exist.

Given that Martin identifies explicitly the concepts or categories which take part in his metaphors, he has no need for further constraints. However, if our concern is to evaluate any expression that comes along for literal or metaphoric coherence, we do have to consider them, and therefore focus on the more general linguistic/semantic question of how verbal concepts and their objects go together. In basing its comprehension on a hypothesized semantic competence which abstracts from literal constraints, then, MAP’s approach is more concerned with *why* we understand--or do not understand--an expression as metaphoric.

CONCLUSION

This discussion of metaphoric coherence has focused on a consideration of how we as language users might extend properties comprehensibly across conceptually different domains. It has been a “fishbowl” presentation in that some of the ambiguities in making coherence and interpretation judgments are readily apparent. However, revealing some of the “hard parts” of metaphor comprehension through a look at conceptually different types of concepts has to be faced sooner or later, if progress is to be made on this issue. The approach presented here suffers from the demand that all lexical items expressing verbal concepts need to be analyzed “abstractly” in terms of a general metaphoric framework, seemingly a quixotic task. However, a careful consideration of what all encountered concepts mean *is* a necessary task if we are to represent human linguistic understanding and the human capability of abstraction. While MAP’s structure-domain system has no claimed psychological validity, it organizes concepts and their extensions and constraints for the purpose of facilitating this task and perhaps pointing out some complexities and limits.

The above observations on constraint criteria for determining cross-modal and other metaphoric relations may give some indication of what we can expect and not expect extensible-property-based representations of constraints to do. Constraints as suggested above can aid in distinguishing metaphoric from other nonliteral or anomalous expressions, and can allow adjustment of an interpretation of “bad” metaphor. However, while we can place some of the burden of ambiguity resolution on discursive context, we have to recognize the relative uncertainty in defining extensible properties, as well as constraints in terms of those properties. This is perhaps a reflection of our inability to completely explain our perception of similarity across conceptual domains.

Computational experience with a more developed system of abstract descriptors, as well as extensions of experiments of the type run by Gentner and France as referred to earlier, may show whether such descriptors are relevant. It could be that reifications of attributes and of other concepts which are not literally “things” would be more effectively described, and their use constrained, in ways other than by features borrowed from literal language. Regardless of the chosen form of constraints, an exploration of why a metaphoric expression is coherent requires a confrontation with semantic abstraction and eventually the fuzziness and subtleties it brings with it.

ACKNOWLEDGMENTS

The author thanks Dan Fass and several anonymous reviewers for their helpful comments on various versions of this paper.

REFERENCES

- Aarts, J. and Calbert, J. 1979. Metaphor and non-metaphor. The semantics of adjective-noun combinations. Max Niemeyer Verlag, Tuebingen, Germany.
- Barnden, J. 1990. Naive metaphysics: A metaphor-based approach to propositional attitude representation. Memoranda in computer and cognitive science, No. M CCS-90-174, Computing Research Laboratory. New Mexico State University, Las Cruces, NM.
- Barnden, J. 1991. Some interplay between metaphor and propositional attitudes. Proceedings of the IJCAI Workshop on Non-Literal Language: Metaphor, metonymy, idiom, speech acts, implicature, Sydney, Australia, pp. 1-11.
- Brugman, C. and Lakoff, G. 1988. Cognitive topology and lexical networks. *In Lexical ambiguity resolution. Edited by S. Small, G. Cottrell and M. Tanenhaus.* Morgan Kaufmann Publishers, Inc., San Mateo, CA.
- Carbonell, J. 1982. Metaphor: An inescapable phenomenon in natural-language comprehension. *In Strategies for natural language processing. Edited by W. Lehnert and M. Ringle.* Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.
- Carbonell, J. and Minton, S. 1983. Metaphor and common-sense reasoning. Technical Report No. CMU-CS-83-110. Carnegie-Mellon University, Pittsburgh, PA.
- Drange, T. 1966. Type crossings: Sentential meaninglessness in the border area of linguistics and philosophy. Mouton & Co., Publishers, The Hague, The Netherlands.
- Fass, D. 1988. An account of coherence, semantic relations, metonymy, and lexical ambiguity resolution. *In Lexical ambiguity resolution. Edited by S. Small, G. Cottrell and M. Tanenhaus.* Morgan Kaufmann Publishers, Inc., San Mateo, CA.
- Fass, D. 1989. Met*: A method for discriminating metonymy and metaphor by computer. Technical Report CSS/LCCR TR 89-15. Simon Fraser University, Burnaby, BC.
- Fass, D. 1991. Metonymy, case role substitution and sense ambiguity. Proceedings of the IJCAI Workshop on Computational Approaches to Non-Literal Language: Metaphor, Metonymy, Idiom, Speech Acts, Implicature, pp. 42-51.
- Fass, D. and Wilks, Y. 1983. Preference semantics, ill-formedness, and metaphor. *American journal of computational linguistics*, 9:178-187.
- Gentner, D. and France, I. 1988. The verb mutability effect: Studies of the combinatorial semantics of nouns and verbs. *In Lexical ambiguity resolution. Edited by S. Small, G. Cottrell and M. Tanenhaus.* Morgan Kaufmann Publishers, Inc., San Mateo, CA.
- Hayes, P. 1985. The second naive physics manifesto. *In Readings in knowledge representation. Edited by R. Brachman and H. Levesque.* Morgan Kaufmann Publishers, Inc., Los Altos, CA.
- Hobbs, J. 1977. Resolving pronoun references. *In Readings in natural language processing. Edited by B. Grosz, K. Sparck Jones and B. Webber.* Morgan Kaufmann Publishers, Inc., Los Altos, CA.
- Hobbs, J., Croft, W., Davies, T., Edwards, D. and Laws, K. 1987. Commonsense metaphysics and lexical semantics. *Computational linguistics*, 13:241-250.
- Indurkha, B. 1986. Constrained semantic transference: A formal theory of metaphors. *Synthese*, 68:515-551.
- Jackendoff, R. 1976. An explanatory semantic representation. *Linguistic Inquiry*, 7:89-150.
- Lakoff, G. 1986. A figure of thought. *Metaphor and symbolic activity*, 1:215-225.
- Lakoff, G. and Johnson, M. 1980. *Metaphors we live by.* Chicago University Press, Chicago, IL.
- Levin, S. 1977. *The semantics of metaphor.* The Johns Hopkins University Press, Baltimore, MD.
- Martin, J. 1990. A computational model of metaphor interpretation. Academic Press, New York, NY.
- Martin, J. 1991. MetaBank: A knowledge-base of metaphoric language conventions. Proceedings of the IJCAI Workshop on Computational Approaches to Non-Literal Language: Metaphor, metonymy, idiom, speech acts, implicature, Sydney, Australia, pp.123-131.
- Norman, D. and Rumelhart, D. 1975. *Explorations in cognition.* Freeman, San Francisco, CA.
- Ortony, A. 1980. Some psycholinguistic aspects of metaphor. *In Cognition and Figurative Language. Edited by R. Honeck and R. Hoffman.* Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.
- Osgood, C. 1980. The cognitive dynamics of synesthesia and metaphor. *In Cognition and figurative*

- language. *Edited by R. Honeck & R. Hoffman.* Lawrence Erlbaum Associates, Inc., Hillsdale, NJ. pp. 203-238.
- Perrault, C. and Allen, J. 1980. A plan-based analysis of indirect speech acts. *American journal of computational linguistics*, 6:167-182.
- Russell, S. Weber. 1975. Computer understanding of conceptually complex phrases. Ph.D. Thesis, Stanford University.
- Russell, S. Weber. 1976. Computer understanding of metaphorically used verbs. *American journal of computational linguistics*, Microfiche 44.
- Russell, S. Weber. 1985. Conceptual analysis of partial metaphor. *In Progress in Artificial Intelligence. Edited by L. Steels and J. Campbell.* Ellis Horwood, Chichester, England.
- Russell, S. Weber. 1986. Information and experience in metaphor: A perspective from computer analysis. *Metaphor and symbolic activity*, 1:227-270.
- Russell, S. Weber. 1988. Metaphor and computational limits. Presented at the Symposium on the Human Dimension of Artificial Intelligence, Louisville, KY. Technical Report No. 89-57, Computer Science Department. University of New Hampshire, Durham, NH. 1989.
- Russell, S. Weber. 1989. Verbal concepts as abstract structures: The most basic conceptual metaphor? *Metaphor and symbolic activity*, 4:55-60.
- Schank, R. 1975. *Conceptual information processing.* North Holland, Amsterdam.
- Schank, R. and Abelson, R. 1977. *Scripts, plans, goals and understanding.* Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.
- Smith, E., Rips, L. and Shoben, E. 1974. Semantic memory and psychological semantics. *In The Psychology of learning and motivation*, 8. *Edited by G. Bower.*
- Suwa, M. and Motoda, H. 1991. Learning metaphorical relationships between concepts based on semantic representation using abstract primitives. *Proceedings of the IJCAI Workshop on Computational Approaches to Non-Literal Language: Metaphor, metonymy, idiom, speech acts, implicature*, Sydney, Australia, pp.123-131.
- Talmy, L. 1978. The relation of grammar to cognition--a synopsis. *In Proceedings of TINLAP-2 (Theoretical Issues in Natural Language Processing. Edited by D. Waltz.* Association for Computing Machinery, New York, NY.
- Talmy, L. 1987. The relation of grammar to cognition. *In Topics in cognitive linguistics. Edited by B. Rudzka-Ostyn.* John Benjamins Publishing, Amsterdam, The Netherlands.
- Thomas, O. 1969. *Metaphor and related subjects.* Random House, New York, NY.
- Tourangeau, R. and Sternberg, R. 1982. Understanding and appreciating metaphors. *Cognition*, 11:203-244.
- Wilks, Y. 1975. A preferential pattern-seeking semantics for natural language inference. *Artificial Intelligence*, 6:53-74.