

John A. Barnden, Mark G. Lee

An Artificial Intelligence Approach to Metaphor Understanding

Abstract

An implemented system called ATT-Meta is sketched. The system can perform an important type of metaphor-based reasoning. It is based on an emphasis on source-domain reasoning and a de-emphasis of the idea of creating new mappings between source domain and target domain. The metaphor-based reasoning is fully integrated into a general framework for uncertain reasoning. The system thereby copes with various different types of uncertainty involved in metaphor understanding. The view of metaphor on which the system is founded is highly understander-relative, and does not require prior assumptions about what literal discourse is.

1. Introduction

Relatively few major metaphor-processing systems have been implemented. (See Martin 1996 for a review. Probably the main systems, apart from our own, are those of Fass 1997, Martin 1990, Veale & Keane 1994, and Narayanan 1997.) But even those that have leave many questions unanswered. In particular, we wish to fill a gap concerning the reasoning that is needed in metaphor understanding: a lack of appreciation of the importance of source-domain reasoning. Source-domain reasoning is reasoning conducted within the terms of the source domain (vehicle domain). We claim that many metaphorical utterances rely heavily on source-domain reasoning for their understanding. Some authors have incorporated source-domain reasoning (notably Hobbs 1990, Narayanan 1997), but have either not implemented a metaphor system (e.g., Hobbs) or have not provided as powerful an implemented framework as that in our implemented system, ATT-Meta.

Previous work on metaphor mostly focuses on processing of the following sorts, in some combination:

- (a) discovering an analogical, structural mapping between the source domain and the target domain, and
- (b) creating new target-domain structures that parallel existing source-domain structures, via whatever analogical mapping has been established.

These can be important processes, but undue attention has been paid to (a), partly as a result of an over-concentration on totally novel metaphor (from the understander's point of view) - that is, metaphorical utterances where the understander does not yet have a mapping between the source domain and target domain. Naturally, in such cases (a) is needed. However, our own focus is on cases where the metaphorical utterance involves a mapping that the understander is already familiar with—so we can say that the utterance rests on a familiar metaphorical view—but where the utterance goes beyond the known mapping- so we can say that it is view-transcending. An example of a view-transcending utterance could be the following sentence:

- (1) “The two ideas were in different store-rooms in John’s mind”.

Suppose that the understander is familiar with viewing minds as physical spaces or containers but possesses no mapping for rooms within a mind-space. The notion of store-rooms in John’s mind is therefore view-transcending. It is the task of source domain reasoning to link view-transcending notions in a metaphorical utterance to the existing mapping. In dealing with an utterance resting on and perhaps transcending a familiar mapping, we adopt a default that goes in the opposite direction to (a):

Map-Extension Minimization: by default, new source/target correspondences to supplement the ones in the existing mapping should not be sought.

New correspondences should only be sought when there is a specific need for them. This principle contributes significantly to computational straightforwardness and efficiency in metaphor understanding, and may contribute also to psychological plausibility. The principle also appears operative in the work of Hobbs (1990) and Narayanan (1997), whereas Martin (1990) does assume that new correspondences should be sought.

A far-reaching characteristic of our approach is its stress on the important role that *uncertainty*, of various different types, plays in metaphor understanding. Uncertainty is recognized in the metaphor area as being important, but the different issues have not previously been teased apart, and few developers of detailed schemes have incorporated anything like an adequate treatment; and even those that have worked towards such a treatment have not explicitly recognized that

uncertainty deeply affects some major qualitative issues in metaphor understanding. An example is the question of how to handle conflicts between target-domain information and transfers from the source domain, and another is the possibility of conflicts between different mappings within a single metaphorical view.

Another contrast to most other work on metaphor is that we fully and explicitly embrace the notion that metaphor is a highly *understander-relative* matter (but see Goatly 1997: p.110 for a comment favourable to understand- relativity). The questions of whether an utterance is metaphorical, in what way it is metaphorical, and of what metaphor(s) it manifests are treated as understand- relative matters, not ones that can be decided without reference to specific understanders. They are relative at least to the understander's lexicon and the understander's familiarity with particular metaphorical views.

Our approach is implemented, in advanced prototype form, in a Prolog program called ATT-Meta. ATT-Meta respects the points listed above to varying degrees - we do not claim a perfect or full approach on all these fronts. ATT- Meta is currently only a reasoning system—the user supplies to it hand- constructed logical semantic forms that could hypothetically be derived from superficial semantic processing of natural language inputs. The construction of a front-end for dealing with text directly is an important long-term aim of our project, but our focus so far has been on the reasoning necessary for handling metaphor in language. This reasoning is a very difficult matter especially because of the needed uncertainty and the connotative richness of metaphor.

The rest of the paper is structured as follows. Section 2 explains some of our basic notions concerning metaphor. Section 3 explains a major concern of this paper, namely the centrality of uncertainty, of various different types, in metaphor. Section 4 describes ATT Meta's metaphor-based reasoning facilities, largely through the medium of one quite complex example on which we have run the system. Section 5 explains how ATT-Meta copes with the types of uncertainty presented in section 3. Section 6 concludes.

2. Some Basic Notions

We give here our own rough working definitions of metaphor and related notions.

2.1. Metaphor and Metaphorical Views

We take *metaphor* in general to be the phenomenon of representing, making statements about, clarifying, querying or otherwise “discussing” aspects of some topic or domain by representing them as if they were, or representing them as similar to, aspects of

a different topic or domain. The first topic/domain is the “target” and the second is the “source.” The representation of target aspects as (or as similar to) source aspects is by means of representational resources that can be used for discussing the source. The “discussing” of the target is done via a mapping between the topics. The mapping specifies that certain aspects of the source correspond to certain aspects of the target, and typically takes the form of a structured analogy between the topics (see notably Falkenhainer et al. 1989). The aspects of target and source can be individual entities, properties, relationships, categories, propositions, etc.

We call the mapping involved here a *metaphorical* mapping or view. As an example, a metaphorical view of ideas as physical objects could include a correspondence between certain ideas and some (hypothetical) physical objects, and the mapping of the notion of physical interaction of such physical objects to the notion of mental activity involving the corresponding ideas.

Metaphorical views correspond to the “conceptual metaphors” of Lakoff (1993). A metaphorical view is a conceptual item, not something that intrinsically involves any specific representational medium such as language. A metaphorical view is either one the understander U already possesses in its entirety, in which case the view is *familiar* to U, or one that U creates, at least in part, during the attempt to understand the utterance itself; in which case the view is (in part) *novel* to U.

2.2. Connotations, Metaphorical Utterances and Manifestations

A *connotation* of an utterance (metaphorical or not) is simply some proposition, question, etc. concerning some topic(s) that is drawn by some semantic/pragmatic processing from the utterance. We usually avoid the word “meaning” in favour of the term connotation.

An utterance is *metaphorical*, for a given understander U, if it engages in metaphor, relative to U. That is, U takes the utterance to imply at least one connotation C about some topic (the target), but some of the surface-linguistic elements used in the utterance have semantic import for the understander in another domain (the source), such that U can, at least in principle, derive C from

- that import
- the import (in whatever domain) of the rest of the utterance
- some mapping that U assigns between the source and target
- and possibly some inferencing in the terms of the target and/or source.

The surface-linguistic elements can be morphemes, words, phrases, phrasal templates, linguistic constructions, etc. They play the role of the representational resources mentioned in our characterisation of metaphor. The mapping is of course a metaphorical view, and we say that the utterance is *manifestation* of that view for U.

For example, consider an utterance of “McEnroe killed Connors,” taken by U to have the connotation C that McEnroe defeated Connors in a tennis match. The word “kill” has the semantic import of causing death in the PHYSICAL COMBAT domain, which is the source. If there is a mapping from PHYSICAL COMBAT to sports that includes a correspondence between dying to losing in a game, then U can work C out. So the utterance is a manifestation of the view of tennis as physical combat.

Notice that U is not necessarily obliged to work out connotations by actually going via the source domain and the mapping, even if U could in principle do so. For example, in the McEnroe/Connors example the verb “kill” might be in U’s mental lexicon together with a meaning such as “defeat [at a sport],” as well as being listed with a cause-death meaning. Then U could derive the connotation C directly. This is so even if U also possesses the means to work out C via the PHYSICAL COMBAT meaning and the above mapping. In that case the utterance is a manifestation of the metaphorical view in question for U, despite the presence of the sports-defeat lexicon sense for “kill”

3. Uncertainty in Metaphorical Reasoning

Uncertainty arises in the processing of metaphorical utterances in a variety of ways, over and above the normal types of uncertainty arising from ordinary lexical ambiguity, syntactic ambiguity, anaphora resolution, and so on. The present section briefly surveys the types of uncertainty involved, concentrating on the case of utterances based on familiar metaphorical views but possibly transcending those views.

3.1. Uncertainty about Relevant Transfer

The main type of uncertainty that has been studied in the metaphor area is uncertainty as to what aspects of the source domain to transfer (in modified form because of source to-target mapping relationships) to become new assertions about the target.

3.2. Uncertainty Concerning the Metaphorical View

It is often unclear which, if any, familiar metaphorical views are manifested by an utterance. For example, in the store-rooms example it is clear that the mind is being viewed as being (or at least containing) some sort of physical space or

container. But what about the ideas, specified as having locations in this space? The understander may be familiar with viewing ideas as physical objects—but is that view certainly needed in the current case? Protoypically, entities stated as having a location in a store-room are physical objects. However, events (at least physical ones) can also be stated to have locations. So we can say that “There was a lot of crying in the cinema that night.” Therefore, it cannot be certain that the ideas are being viewed as physical objects as opposed to physical events (or possibly other sorts of entity).

3.3. Uncertainty within the Source

The knowledge the understander has about the source domain, at least if it is a commonsense domain, is likely to be full of uncertainty. For example, in understanding the store-room sentence, the understander may need to use knowledge that physical objects that are in different store-rooms in a building are *presumably* not close together (relative to the overall size of the building). If there is a mapping from physical closeness to ability to interact mentally, inferring the lack of closeness could be important. But in unusual circumstances, the objects could be close together. So, the inference drawn in the source domain must be uncertain. Hence, the target-domain connotation to be drawn can at its strongest be that the ideas in question are *presumably* not able to interact mentally.

3.4. Conflicts between Source and Target

Clearly, also, uncertain reasoning within the target domain may be needed. The target information resulting directly from applying mapping relationships in the metaphorical view may cooperate or conflict with other information in the target domain. We concentrate on the case of conflict here.

It is commonly assumed, at least tacitly, that in cases of conflict between the metaphorically-derived information and existing, non-metaphorically-derived information about the target, the latter should win. However, this is only self-evident when the latter is certain (and the former is uncertain—which it always is). When one allows for uncertainty of information in the target domain, there is no particular reason why that information should win. Indeed, the point of a metaphorical utterance may precisely be to depict an exceptional situation that contradicts an existing target default. For example, if one says “those arrangements were in totally different parts of my mind” as an excuse for not noticing that one has double-booked oneself, the utterance is providing an exception to a default that related ideas about daily arrangements should normally interact mentally.

4. Metaphor-Based Reasoning in ATT-Meta

The current metaphor-based reasoning facilities in ATT-Meta are designed to cope with view-transcending manifestations of familiar metaphorical views. ATT-Meta copes, to some promising initial extent at least, with all the types of uncertainty listed in section 3. Part of the purpose of the present section is to suggest by example that it is reasonable to claim that the Map-Extension Minimization default stated in the Introduction is appropriate: source-domain aspects of a metaphorical utterance that are not mapped by existing mappings should (by default) not be given mappings to the target domain. We only sketch ATT-Meta's operation in gross outline. Additional information on ATT-Meta's representations, processing and reasoning abilities can be found in Barnden (1998, Barnden & Lee (1999) and elsewhere.

Consider again the store-room sentence:

“The two ideas were in different store-rooms in John's mind”.

The sentence manifests two metaphorical views: MIND AS PHYSICAL SPACE and (probably) IDEAS AS PHYSICAL OBJECTS. ATT-Meta is familiar with both metaphorical views. But ATT-Meta only has a physical sense for “storeroom,” and the notion of physical store-room is not mapped by either of the metaphorical views to any target-domain notion.

We try to minimize the activities of discovering or inventing mappings for view-transcending notions such as store-room. Rather, we rely on *on-the-fly source-based inference-inference* conducted, during understanding, within the terms of the source domain(s) of the metaphorical view(s) involved—to link the utterance's use of the notions to notions that *are* mapped by the views. By this means, the understander can infer connotations of the utterance such as the following:

Connotation

The mentioned ideas were involved in John's mind in such a way that John was NOT in a position to mentally operate upon them conjointly—for instance, to compare them or to perform an inference that directly relied on both of them.

We proceed to explain, with the aid of Fig. 1, how the connotation can be derived, as an uncertain inference obtained by usage of

- A. the information in the utterance *taken in terms of the source domain*, i.e. taken to say that John's mind has physical store-rooms as parts and that the ideas are in those store-rooms:

- B. general knowledge about real physical store-rooms and other physical objects, locations and interactions;
- C. *conversion rules* associated with MIND AS PHYSICAL SPACE and IDEAS AS PHYSICAL OBJECTS, where a conversion rule is a type of context-bridging rule that maps information between the source domain and the target domain of a metaphorical view

The reasoning, apart from the application of the conversion rules, occurs within a special context that we call a *source-based pretence cocoon*. Within this cocoon, the *source-domain* information from the sentence—(A) above—is taken to be true, no matter how absurd. ATT-Meta can use any of its reasoning rules within the source-based pretence cocoon to try to infer propositions that can be used by the conversion rules to generate propositions about the target domain. The propositions thus generated are in the system’s top reasoning context, not the pretence cocoon. The reasoning within the pretence cocoon is what we called above the on-the-fly source-based reasoning. Because of the nature of the conversion rules and information (A), this reasoning is largely or wholly conducted using rules about the source domains. (But in fact we allow it to involve other rules also, such as rules about the target-domain itself.)

The main conversion rule used maps inability of someone to physically operate on two ideas conjointly (where the ideas are viewed metaphorically as a physical object) to that person’s inability to operate *mentally* on the two ideas conjointly. (We emphasize that the word “operate” here is used loosely in both its physical and mental sense. Physical operation includes perception, comparison, etc; mental operation similarly includes noticing, comparison, etc.) The rule is actually as follows:

(2)

IF J is an idea AND K is an idea

AND it is being pretended that, presumably at least, J is a physical-object

AND it is being pretended that, presumably at least, K is a physical-object

AND X is a person

AND it is being pretended that, presumably at least, NOT(X can-physically- operate-on {J,K})

THEN [presumably] NOT(X can-mentally-operate-on {J,K}).

A hypothesis glossed here as “it is being pretended that, presumably at least, H” for some hypothesis H is a statement that the pretence cocoon contains H and that the level of certainty for H within that cocoon is at least presumed. This is one of the four positive certainty levels in ATT-Meta, the other three being certain, suggested and possible. The latter two indicate weaker levels of certainty

than presumed (see Barnden 1998 or Barnden & Lee 1999 for an explanation). When a hypothesis is tagged as presumed, it is regarded as a working hypothesis (i.e., a default conclusion that can potentially be upset by further evidence).

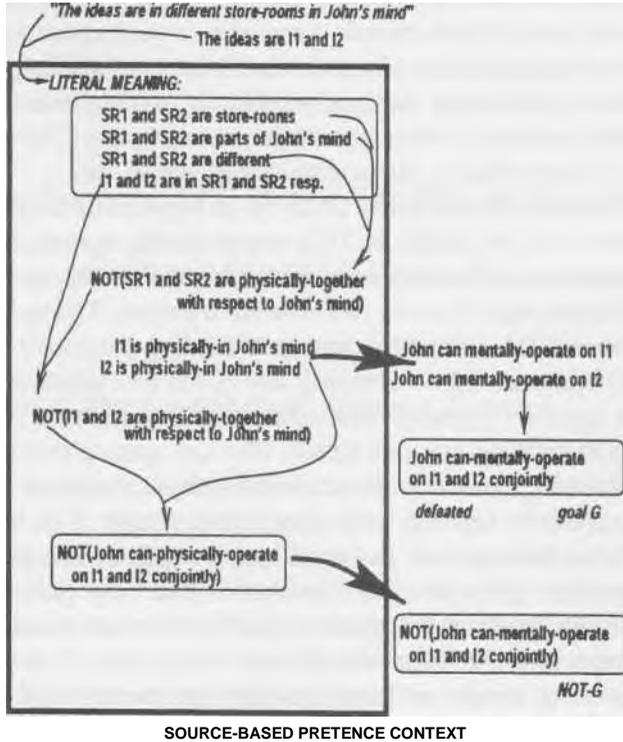


Fig. 1. Part of the reasoning in store-room example. The source-domain meaning of the sentence is shown in the topmost part of the large box. The processing shown within the large box is source-based inference. The processing outside the box is reasoning within the terms of the target domain. The arrows from within the box to outside depict applications of conversion rules.

A hypothesis of form "it is being pretended that, with certainty at least 1, H" can be expressed as a formula of the form pretend (H, 1) . Such a formula is outside the pretence cocoon. The formula H is within the cocoon. A formula using the pretend operator is called a "pretence hypothesis". To oversimplify this, every such hypothesis is reflected within the cocoon by its "H" argument, and every hypothesis H within the cocoon is reflected outside by a pretence hypothesis.

The "presumably" just after the THEN in the rule above acts as a limit on the certainty that the rule can produce for its conclusion. For example, even if

all the conditions of the rule are matched by absolutely certain hypotheses, the conclusion will only be presumed.

Conjointness of operation is modelled in (2) by having a set of objects, rather than a single object, as an argument in the applications of *can-physically-operate-on* and *can-mentally-operate-on*. The above conversion rule is one of a set of related rules. Others include a converse, a contrapositive and a conversecontrapositive for the above. (Two of these map from the target domain to the source domain, rather than the reverse. This is another unusual feature of our approach with respect to other approaches to metaphor.) Also, there are rules dealing with single objects, rather than non-singleton sets.

The connotation shown above arises by an application of (2). This is because it can be shown that, by default, NOT(*X can-physically-operate-on* {11,12}), where 11 and 12 are the ideas mentioned in the sentence, from the premise that II and 12 are in different physical store-rooms in John's mind. This is because (a) storerooms are rooms, (b) different rooms in a building are usually do not spatially overlap, (c) objects in non-overlapping subregions of a physical region R are, by default, not spatially-together with respect to the scale of R, and (d) in the common-sense physical world, a person who can operate in a physical region R cannot, by default, physically operate on two physical objects in R conjointly if they are not spatially together with respect to the scale of R. In our example, R is instantiated as John's mind. But note that (a) to (d) are just pieces of common-sense information about the physical world. The only properly metaphorical processing is the application of rule (2). Principles (a) to (d) are couched as ATT-Meta rules that we do not detail here.

The reasoning using (a-d) occurs within the source-based pretence cocoon. Into that cocoon is inserted (at the appropriate moment in the backwards rule-based reasoning) the facts that II is in a store-room SRI, 12 is in a store-room SR2, SRI is part of John's mind, and SR2 is part of John's mind. From these facts, ATT-Meta can infer, within the source-based pretence cocoon, that John's mind is a building and therefore a physical region, and II and 12 are physical objects. These inferences use further common-sense rules about the ordinary physical world.

Note that the above reasoning requires it to be shown that John can physically operate within his mind. This is shown by a rule associated with MIND AS PHYSICAL SPACE that says that if a person X's mind is pretended to be a physical region then X can physically operate within it.

The key point is that this reasoning from the *source-domain* meaning of the utterance, conducted within the pretence cocoon, link up with the knowledge displayed as (2). That knowledge is itself of a very fundamental, general nature, and does not, for instance, rely on the notion of store-room. Any line of within-pretence inference that linked up with that knowledge could lead to a conclusion that the person in question could not mentally operate on some ideas conjointly.

Recall that ATT-Meta applies rules in a backwards, goal-directed way. So, we need a goal to start the process off. For the moment, let us without explanation assume that the goal is

(G) John can-mentally-operate-on {II, 12}.

In fact, the system will conclude with a certainty of presumed for the *negation* of (G), thereby drawing the Connotation displayed above.

5. Uncertainty in ATT-Meta's Handling of Metaphor

The way ATT-Meta cope with the types of uncertainty listed in section 3 can be illustrated as follows.

5.1. Uncertainty about Relevant Transfer

In the ATT-Meta reasoning sketched above, we just assumed without explanation that the question to be investigated is whether the ideas interact mentally. In a full discourse understanding system, however, this question would normally be apparent from the context. For instance, the surrounding discourse might be about some proposition that John fails to infer, where the proposition follows from the two ideas if taken conjointly, but that does not follow from either of them separately. The store-room sentence could be offered as an indication of why John did not infer the proposition.

But if context does not raise the issue of the two ideas interacting mentally, then (other things being equal) ATT-Meta would not be led to perform the conversion- rule application that is shown in Figure 1. We propose that the connotations extracted from a metaphorical utterance are extracted in the service of trying to find coherence with surrounding discourse. The details of the context affect precisely which connotations are relevant. Notice that in the store-room example, a connotation we have not discussed is that each idea individually takes a background role in the agent's overall thinking. It might be this connotation that links to the surrounding discourse, and not the connotation that the ideas do not interact mentally.

5.2. Uncertainty Concerning the Metaphorical View

ATT-Meta may merely have 'presumed', for instance, as a (tentative) level of certainty for a pretence hypothesis, such as the hypothesis that it is pretended

that idea II is a physical object, in the store-room example. This hypothesis is then potentially subject to defeat. Indeed, note that, in the example the pretence hypothesis just alluded to only arises in the middle of reasoning: the system does not start with a premise that some idea is being viewed as physical object. Equally, the system does not start with a premise that John's mind is being viewed as a building.

Note that the uncertainty of the pretence hypothesis (the hypothesis that it is pretended that idea II is a physical object) is a separate matter from the uncertainty of the within-pretence hypothesis that II is a physical object. That hypothesis happens also only to be presumed rather than certain.

5.3. Uncertainty within the Source

In ATT-Meta the hypotheses and reasoning within the pretence cocoon are usually uncertain. For instance, U is not *certain* that a person cannot physical operate conjointly on two objects if they are physically separated from each other. There could be evidence in a particular case that the person can indeed operate conjointly on them.

5.4. Conflicts between Source and Target

Conversion rules like (2) are merely default rules. There can be evidence against the conclusion of the rule. Whether the conclusion survives as a default (presumed) hypothesis depends on the relative strength of the evidence for and against the conclusion. Thus, whether a piece of metaphorical reasoning overrides or is overridden by other lines of reasoning about the target is matter of the peculiarities of the case at hand. The evidence on either side is weighed by a specificity-based scheme that is not tailored for metaphor, so non- metaphorically-based information has no special advantage.

6. Conclusion

ATT-Meta implements a type of metaphor-based reasoning, applicable to a major class of metaphorical utterances in discourse. This class is typified by the storeroom example (1), where view-transcending aspects of the source domain of a familiar metaphorical view are used in such a way that there is no reason to think that these aspects need to be mapped to the target domain. ATT-Meta tries to exploit existing mappings as much as possible through source-domain reasoning

that links unmapped source-domain aspects of the utterance to mapped source-domain notions. Indeed, ATT-Meta cannot at present construct new mappings at all. The determination of when such mapping construction should occur, and how it should be done, are questions for future research.

ATT-Meta joins metaphor-based reasoning and uncertainty-handling in a unified framework. It can cope with major types of uncertainty associated with metaphor, at least to a promising extent.

We do not claim that ATT-Meta can deal with all the subtleties of the types of metaphorical utterance it is directed at. In particular, ATT-Meta currently lacks a proper treatment of change and time, and so cannot do justice to the processual quality of many metaphorical descriptions. But one issue partially addressed is the context-sensitivity of what particular connotations are drawn, because the metaphorical reasoning (as with all other reasoning in the system) is goal-driven and therefore responsive to the particular questions raised by surrounding sentences.

We have been careful to provide characterizations of such notions as metaphor and manifestations of metaphors in a way that makes no reference to the notion of “literal” (which is fortunate in view of the notorious slipperiness of this term - see, e.g., Lakoff 1986). Instead of appealing to literal meanings of words, we have appealed to the more straightforward notion of source-domain meanings. The source-domain meanings could themselves be the result of figurative processes. The appeal to source-domain meanings as opposed to literal meanings is strongly correlated with our understander-relative view of metaphor, because the question of what domains it is that a given word has lexicon-provided meanings in is beneficially viewed as a highly understander-relative matter.

Acknowledgment

The research was supported in part by grant number IRI-9101354 from the National Science Foundation (USA) and grant number GR/M64208 from the Engineering and Physical Sciences Research Council (UK).

References

- Bamden, J. A. (1998). Combining *uncertain belief reasoning and uncertain metaphor-based reasoning*. In Procs. Twentieth Annual Meeting of the Cognitive Science Society, pp. 114-119. Mahwah, N. J.: Lawrence Erlbaum Associates.
- Bamden, J. A. & Lee, M. G. (1999). *An implemented context system that combines belief reasoning, metaphor-based reasoning and uncertainty handling*. In P. Bouquet,

- P. Brezillon & L. Serafini (Eds), *Second International and Interdisciplinary Conference on Modeling and Using Context* (CONTEXT'99, Trento, Sep 9-11, 1999), Lecture Notes in Artificial Intelligence, 1688, pp. 28-41, Springer.
- Falkenhainer, B., Forbus, K. D. & Gentner, D. (1989). *The Structure-Mapping Engine: algorithm and examples*. Artificial Intelligence, 41 (1), 1-63.
- Fass, D. (1997). *Processing metaphor and metonymy*. Greenwich, Connecticut: Ablex.
- Goatly, A. (1997). *The language of metaphors*. London and New York: Routledge.
- Hobbs, J. R. (1990). *Literature and cognition*. CSLI Lecture Notes, No. 21, Center for the Study of Language and Information, Stanford University.
- Lakoff, G. (1986). *The meanings of literal*. Metaphor and Symbolic Activity, 1 (4), pp. 29-36.
- Lakoff, G. (1993). *The contemporary theory of metaphor*. In A. Ortony (Ed.), *Metaphor and Thought*, 2nd edition, pp. 202-251. Cambridge, U.K.: Cambridge University Press.
- Martin, J. H. (1990). *A computational model of metaphor interpretation*. Academic Press.
- Martin, J. H. (1996). *Computational approaches to figurative language*. Metaphor and Symbolic Activity, 11, pp. 85-100.
- Narayanan, S. (1997). *KARMA: Knowledge-based action representations for metaphor and aspect*. Ph.D. thesis, EECS Department, U. of California, Berkeley, August 1997.
- Veale, T. & Keane, M. T. (1994). *Belief modelling, intentionality and perlocution in metaphor comprehension*. In *Procs. Sixteenth Annual Conference of the Cognitive Science Society*, pp. 910-915. Hillsdale, N. J.: Lawrence Erlbaum.