

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The Utility of Reserved Transfers in Metaphor

Permalink

<https://escholarship.org/uc/item/3vw4s91k>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 23(23)

ISSN

1069-7977

Author

Barnden, John A.

Publication Date

2001

Peer reviewed

The Utility of Reversed Transfers in Metaphor

John A. Barnden (J.A.Barnden@cs.bham.ac.uk)

School of Computer Science; University of Birmingham
Birmingham, B15 2TT, United Kingdom

Abstract

In metaphor research there is usually some notion of transfer of aspects of the source domain to the target domain. More rarely, transfers in the opposite direction are countenanced, affecting one's perception of source as well as target. This paper argues that, even without this aim, transfers from target to source should happen. One radical claim here is that it is often better to translate information from literal sentences into prevailing metaphorical terms than to translate the information from metaphorical sentences into literal terms. The issues have been obscured by confusion between intuitive directions of static source/target mappings, directions of individual transfer actions, and direction of main intended information flow. Relevance to an implemented AI system for metaphorical reasoning, ATT-Meta, and to Blending Theory are briefly mentioned. Asymmetry of metaphor is also addressed.

Introduction

In this paper, metaphor is thinking or communicating about some target scenario TS in a way that relies on or is motivated by seeing it as something one takes to be qualitatively different from it. Consider: “*In the far reaches of her mind, Anne knew Kyle was having an affair*” (from real discourse: Gross, 1994). TS is what is going on in Anne's mind. Her mind is being seen as a physical space that has “far reaches.” The utterance relies on or is motivated by a metaphorical view of MIND AS PHYSICAL SPACE. In this view, the *source domain* is the knowledge domain concerned with physical spaces, locations, etc., and the *target domain* is concerned with minds and mental states/processes. My term *metaphorical view* means much the same as *conceptual metaphor* (Lakoff, 1993). However, I use a different term partly because Lakoff makes claims about conceptual metaphors that do not affect the present paper.

In analyses of metaphor there is usually a notion of *transfer* of aspects of the source domain to the target domain. The transfer involves copying or in some way translating the source aspects. An “aspect” is an entity, property, relationship, proposition, ... that features in source-domain knowledge. The centrality of *source-to-target* (S→T) transfer is especially evident in feature-transfer accounts (e.g., Ortony, 1979) and analogy-based accounts (e.g., Falkenhainer *et al.*, 1989; Holyoak & Thagard, 1989). In the former, understanding of metaphorical utterances is meant to proceed by

finding one or more suitable features of a source entity and ascribing them to, or emphasizing them in, a target entity. In analogy-based accounts, the understander either already possesses an S→T mapping handling some aspects of the source or constructs such a mapping on the fly. The understander uses that mapping in transferring *further* aspects of the source to become potentially new target aspects, or at least target aspects that were not previously attended to or not dealt with in the original mapping. The rest of the paper will be geared ostensibly to analogy-based accounts. However, observations to be made generalize to other accounts as long as they involve notions of mapping and/or transfer.

A mapping generally maps more than one aspect of a domain. I call a part of the mapping that is concerned with one specific aspect a *mapping relationship*. Consider the Socrates-as-midwife metaphorical view in Plato's *Theaetetus*, analyzed by Holyoak & Thagard (1989). This involves, amongst others, a mapping relationship between a midwife and Socrates, one between students (even if male) and mothers, one between babies and ideas, and one between pregnancy and idea-development.

Mappings must be distinguished carefully from actions related to them. A mapping is just a set of relationships between aspects of two domains. It is to be distinguished from the act or process of creating the mapping (another possible meaning of the word “mapping”). Also, a transfer is an action. In fact, a transfer generally rests in part on acts of using the existing mapping. For instance, we might transfer the source-domain proposition that *Socrates helps Theaetetus to give birth to [a particular idea]* to become the target-domain proposition that *Socrates helps Theaetetus to produce [that idea]*. Here the particular proposition transferred was not previously mapped, but its parts were. There could be more creative transfers as well, such as transferring the cleaning up of the afterbirth to eliminating useless side-effects of a produced idea. Either sort of transfer involves actions exploiting existing mapping relationships (e.g., the one from babies to ideas).

Metaphorical views are generally (and perhaps always) considered to be directed, in a natural and intuitive way. A view of midwives as teachers is clearly distinct from a view of teachers as midwives. This is the asymmetry of metaphor. In a view of A as B, the *direction*

of the metaphor is from B to A. Also, in all mapping-based accounts I have seen, the *intuitive direction of the mapping* is usually taken implicitly or explicitly to be the same as the direction of the metaphor.

But there is a further direction, namely *the direction of information flow*. In most examples in the literature, the information flow is from source to target, in that all that is discussed is informational effects on the target. However, in some accounts, notably interaction accounts (Black, 1979; Waggoner, 1990) and the blending-based account (Turner & Fauconnier, 1995), information flow from target to source (T→S) is allowed for, though much less commonly analyzed than S→T flow. Thus, in general, the direction of the metaphor may not be the same as the direction of all information flow. A major task of this paper is to expand on this point beyond where the literature has taken it so far, showing that it is of much more general importance than heretofore realized.

The plan of rest of the paper is as follows. The next section explains why the direction in which a mapping relationship is used (e.g., during a transfer) is theoretically independent of its intuitive direction. The third section shows argues for T→S as well as S→T information-flow and mapping-usage actions. The fourth section explains that it is inappropriate to think of the direction of mapping always being (wholly) the same as the direction of the metaphor. The fifth section comments on the asymmetry of metaphor. The sixth links the considerations to an implemented AI system, ATT-Meta, that performs some of the reasoning needed in metaphor understanding (Barnden *et al.*, 1996; Barnden & Lee, 1999, 2001).

Directionality: Relationship versus Usage

As stated above, mappings are usually described informally as going from source domain to target domain. (Some authors, e.g., Holyoak & Thagard, 1989, occasionally depart from this). Each mapping relationship can just be denoted as an ordered pair (s,t) where s is the source-domain aspect mapped and t is the target-domain aspect it maps to. For instance, in Socrates-as-midwife, a source-domain scenario containing a particular midwife *mw1* is assumed, and we have the mapping relationships (*mw1*, *Socrates*) and (*give-birth*, *produce-idea*).

Now, it is normally assumed that the *use* of a mapping relationship (s,t) is in the direction from s to t . At one point in processing, a source-domain structure involving s may be being worked on; and then, typically as a result of an attempted transfer, a structure involving t will be considered. But, in principle, the direction of use is independent of the direction of the relationship. If for some reason it were beneficial to create source-domain structures that paralleled existing target-domain ones according to the metaphorical view at hand, a mapping relationship (s,t) could be used, in reverse, to go from t to s . Equally, we could just as well have mapping relationships that intuitively go from target to source without affecting their *usability* from source to target.

In much work on analogy and metaphor, mappings are

required (e.g., in SME, Falkenhainer *et al.* 1989) or preferred (e.g., in ACME, Holyoak & Thagard, 1989) to be one-to-one. A one-to-one mapping does not allow there to be two different mapping relationships $(s,t1)$ and $(s,t2)$ or two different mapping relationships $(s1,t)$ and $(s2,t)$. Clearly, if a mapping violated the former condition, some attempted S→T transfers would be faced with extra complication because of the choice between target aspects; and a mapping violating the latter condition would similarly complicate some attempted movements from target to source. These difficulties do not, however, stop a particular (s,t) or (t,s) relationship in a non-one-to-one mapping being usable in either direction.

Usefulness of Target-to-Source Transfers

Some accounts of metaphor (notably interaction accounts and Blending Theory) allow for T→S transfers. Such transfers therefore use mapping relationships in the T→S direction. However, in those accounts, attention is focused on cases where the ultimate effect is to make some relatively long-term change in the understander's appreciation of the source domain. In contrast, the present paper argues that T→S transfers can be useful even when there is no effect on the source domain that outlives the short-term purposes of the current processing (e.g., understanding a sentence), and where the original goal of the processing is purely to add information to the target domain. The general argument is that T→S transfers can create source-domain information that feeds into within-source-domain processing that in turn ultimately feeds back into some S→T transfer. We therefore have a distinction between *direction of ultimate information flow*, which is normally S→T, from *direction of individual transfers*, which can be T→S, although there must be at least one S→T if ultimate information flow is to be in this direction.

The next three subsections look at different types of T→S transfer.

T→S Transfer: Certainty Downgrading

In a teacher-as-midwife scenario, suppose that reasoning within the source domain establishes, to some level less than absolute certainty, that *Adonis [a student] gave birth to the idea J* (J viewed as a child). Let us call this proposition SP. Suppose that by ordinary S→T transfer this creates the target-domain proposition TP that *Adonis produced J*. But, finally, suppose that there is an independent argument in the target domain that Adonis did not produce J, and that this argument is deemed stronger than the metaphor-based argument. Thus the certainty level ascribed to TP must be downgraded. Now, in discussions of such conflict in the literature, it is not pointed out that *therefore* it may be desirable or perhaps even necessary (a) to downgrade correspondingly the certainty level ascribed to the source-domain proposition (here SP) from which the downgraded target-domain proposition (here TP) came, and possibly also (b) to affirm the negation of SP.

Action (a) is T→S transfer of certainty downgrading. The motivation for it is that the original source-domain proposition (SP) could have been used to support other propositions in the source domain scenario (e.g., that *Socrates has acted as a midwife for Adonis*). The downgrading of SP may therefore be needed so that those other propositions can be properly downgraded—and this may then require withdrawal of earlier S→T transfers of those propositions. Somewhat similarly, affirmation of the negation of SP in the source domain could be useful as it could lead to new inferences in the source domain and hence new S→T transfers.

Metaphorization of the Literal

There is a much stronger and more explicit type of T→S transfer. Consider the following discourse fragment:

Socrates helped Adonis to give birth to [the idea J].
Similarly, John helped Mary to produce [some idea K].

It would surely be natural to take John's help to be metaphorically a matter of helping a birth. That is, to transfer the target-domain proposition that *John helped Mary produce K* to become the source-domain proposition that *John helped Mary give birth to K*. Then, the rich resources of the source-domain scenario are available to make further inferences. One such inference, albeit an uncertain one, could be that *John acted as a midwife for Mary*. There could then be a further inferences that John was instrumental in introducing Mary to the responsible sexual partner. This matchmaking function of (ancient Greek) midwifery is explicit in *Theaetetus*. Such further propositions could then lead by ordinary S→T transfer to new propositions in the target domain, for instance that John introduced Mary to people who stimulated her ideas.

Notice the contrast to the following method: derive target domain propositions from the first, metaphorical, sentence; then extract target-domain propositions from the second, non-metaphorical, sentence; then integrate the two sets of propositions. Most existing accounts of metaphor, in not properly dealing with the role of metaphorical utterances in *discourse* leave one with the impression that this would have to be the method. But it is highly impoverished compared to the method in the preceding paragraph, as it does not access the rich source-domain scenario. The impoverished method does allow room to infer that John is Socratic, partly because of the word "similarly"; but once the metaphor used in the first sentence has been left behind there is no strong impetus to make that inference, rather than simply interpreting the word to be pointing out that John is like Socrates purely in having helped someone produce an idea.

In any case, the argument does not depend on the word "similarly." The alternative second sentence "*Socrates also helped Theaetetus to produce [idea K]*" could again be felicitously be processed by doing a T→S transfer to

get the proposition that *Socrates helped Theaetetus give birth to K*.

The argument somewhat relies on the source domain scenario not having a complete, equally rich and extensive correspondent in the target domain. That is, reasoning in the source domain uses much knowledge about midwives and their role in (ancient Greek) society, where that knowledge is not all mapped to corresponding knowledge in the target domain. That this lack of mapping is common in metaphor is argued further in (Barnden *et al.*, 1996; Barnden & Lee, 2001), but is also linked to common themes in metaphor research such as the relative unparaphrasability of many metaphorical utterances (see, e.g., Waggoner, 1990) and the relative familiarity, richness and accessibility of source domains as opposed to target domains (see, e.g., Lakoff, 1993). Indeed, even if the metaphorical mapping captured all the richness of facts in the source domain, that would still not be enough because methods of reasoning peculiar to the source domain may need to be captured as well. Also, if familiarity of subject matter can affect the facility of people's reasoning even if the pattern of reasoning is kept constant (see, e.g., Johnson-Laird, 1983:29–34), source-domain reasoning stands to have an advantage just from this factor.

Also, it is a mistake to think that if there is an extended use of a metaphor across a stretch of discourse then target-domain information has to be derived from *each* metaphorical patch in the stretch. It may only be necessary to switch to the target-domain once some source-domain *conclusion* has been obtained from within-source-domain reasoning that the stretch stimulates. Thus it may not only be fruitful but also much more economical to metaphorize intervening literal statements than to literalize the metaphorical ones.

In sum, in many cases the proper way to integrate metaphorical discourse elements with non-metaphorical ones is *not* to literalize the metaphorical ones and then do integration (that's the impoverished method above) but rather to metaphorize the literal ones and then integrate. This technique is a radical departure from other research on metaphor, even when discourse-sensitive as in Hobbs (1990). Hobbs does not preclude metaphorization, but he does not appear to have argued for it.

T→S Transfer: Reasoning Queries

One sense in which T→S transfers can occur is through query-directed reasoning (goal-directed reasoning). This is a powerful and important technique in AI generally. In particular, it can be used to focus metaphorical processing fruitfully, and this is especially important given the notorious indeterminacy of the process of extracting information from metaphorical utterances. The ATT-Meta approach places great stress on the technique. In query-directed reasoning, the process starts with a query—a question as to whether something holds. Queries are compared to known propositions and/or used to generate further queries, which, if eventually established, could help provide an answer to the original query; and the in-

vestigation of these queries results in turn from recursive application of the same principle, or by a match to given information. Suppose the discourse at hand is using the Socrates/midwife metaphor, and one reasoning query that has been posted in the target domain (as a result of processing of surrounding discourse) is whether a student *produced a particular idea J*. By virtue of suitable mapping relationships, this query could give rise to the query as to whether the student *gave birth to J*. We can say therefore that the target-domain query has been transferred to become a source-domain query.

Query-directed reasoning of this sort in metaphor or analogy has been advocated by others (e.g., Markman, 1997), and for the particular purposes of the present paper it involves a relatively uninteresting sort of T→S transfer, as it is only useful if it eventually leads to a proposition transfer from source to target. In the student-producing-idea example, the task is to find support for the *student-gives-birth* proposition in the source. This supportedness in the source must be transferred to turn the *student-produces-idea* query in the target into a proposition. Thus, the proposition transfer is S→T even though the query transfer is T→S.

Directionality: Mapping vs Metaphor

Given that the intuitive direction of a metaphor of A-as-B is from B to A, it seems obvious that the intuitive direction of the *mapping* involved in the metaphor should be the same. However, this is simplistic—there are exceptions. They may be quite common, and different mapping relationships in a given metaphorical view may intuitively and theoretically be best viewed as going in different directions.

Metaphorical views of IDEAS AS PHYSICAL OBJECTS and MIND AS PHYSICAL SPACE are used in the sentence “*The idea was hidden behind a door in Mary’s mind.*” Note that whereas the mentioned idea is being viewed as a physical object, so there is indeed a physical object in the source-domain scenario that is mapped to the idea, other implied physical objects in the source-domain scenario, such as the door, should presumably *not* be mapped to ideas. If, therefore, one takes the *property* of being a physical object to map to the property of being an idea, one would have to say that not all applications of *being-physical-object* in the source-domain scenario map to *being-idea* in the target domain. This is an unfortunate gap between the property and its applications, and brings into question the idea that the *being-physical-object* property maps to *being-idea* after all.

One could, of course, say that the *being-physical-object* property does not map at all, and that it is only particular physical objects that map. But that is surely less intuitively appealing than having some sort of mapping of a property. In fact, it is better to take the following, reversed, stance: the property *being-idea* maps to the property *being-physical-object*. This is a natural mapping relationship because it is a reasonable assumption that if one idea in the specific target-domain scenario

at hand is being viewed as a physical object then others are too. Consider for instance the following hypothetical discourse segment:

Many ideas were whizzing around inside Mary’s mind. John’s question made her think up even more.

This again uses both IDEAS AS PHYSICAL OBJECTS and MIND AS PHYSICAL SPACE. It would be natural to assume that the further ideas implied by the second sentence are also physical objects inside Mary’s mind-space (and indeed whizzing around in it). This is another example of metaphorization. Thus, it is plausible that the mapping of *being-idea* to *being-physical-object* has uniform application to all ideas mentioned in the local discourse context.

As for MIND AS PHYSICAL SPACE, the point is yet starker. Consider “*The idea was in the recesses of their minds.*” The minds of all the people concerned are being viewed as physical regions. However, *not* all physical regions in the source-domain scenario are being viewed as minds: in particular, the recesses are not. It is therefore much more natural to think in terms of a mind-to-region mapping relationship rather than the other way round.

We can always restrict a not-uniformly-applicable property (or relationship) in such a way that the resulting property (or relationship) does apply uniformly. However, the restriction may be highly unnatural. For instance, in IDEAS AS PHYSICAL OBJECTS the property *being-physical-object*, restricted to apply only to a physical object *that happens to correspond to a particular mentioned idea* trivially coheres with its applications, with respect to the metaphorical mapping. Of course, the italicized restriction is not itself a restriction framed in terms of the source domain.

Given that a metaphorical mapping typically involves several mapping relationships (e.g., applying to different properties), the possibility arises of having its component relationships go in different directions. Also, a given mapping relationship together with its inverse may both be intuitively part of the mapping. To accommodate these possibilities we can either broaden the notion of mapping to allow relationships to go in different directions, or we can replace the single mapping by two mappings, one consisting of S→T relationships and the other of T→S ones.

Asymmetry of Metaphor

It is frequently pointed out (e.g., Way, 1991) that metaphor is asymmetric (see Introduction above). The present paper might be thought to conflict with this, as it claims that the direction of the metaphor does not completely determine the direction of transfers or even the intuitive direction of mapping relationships. However, there is no conflict, because of the following points.

There must still be some information flow from source to target, i.e. in the direction of the metaphor: it is the *target* domain that contains the topic being attended to. Even if teachers-as-midwives and midwives-as-teachers

both happened to use exactly the same mapping relationships (up to inversion), it is very different to conclude, say, that in reality certain teachers help certain students from concluding, say, that in reality certain midwives help certain pregnant mothers.

But, in any case, the present paper in no way implies that A-as-B would indeed involve exactly the same mapping relationships as B-as-A, even though in practice there may well be considerable overlap, especially as a structural isomorphism between parts of two domains is an inherently symmetric thing. The isomorphism that is appropriate for one direction of metaphor may be slightly or greatly different from the one appropriate to the other direction. (For one thing, even with a fixed direction there can be competing possibilities for partial isomorphism. Such competition is an important aspect of SME and ACME.) Whereas in some particular discourse a use of teachers-as-midwives might involve an isomorphism between the process of giving birth and the process of producing an idea, a use of midwives-as-teachers in another discourse might rest on an isomorphism between the process of a mother coming to bond with her already-born baby and the process of a person producing an idea.

Finally, mapping relationships for A-as-B and B-as-A could be intuitively similar but be different in detail. We saw that for IDEAS AS PHYSICAL OBJECTS, not all physical objects in the source-domain scenario are mapped to ideas, whereas all ideas in the target-domain scenario are likely to be mapped to physical objects. If a metaphorical view of PHYSICAL OBJECTS AS IDEAS were to be used in discourse, it could similarly be that not all ideas in the source-domain scenario were mapped to physical objects but that all physical-objects in the target-domain scenario were mapped to ideas. Asymmetry is addressed again at the end of the next section.

The ATT-Meta System

The ATT-Meta system is too complex to be described at any length here. It is described in Barnden & Lee (1999, 2001). The present section summarizes how the system is related to the issues in previous sections.

The system is aimed at performing the reasoning needed for understanding a broad class of metaphorical utterances that we call *map-transcending utterances based on familiar metaphors*. Here the understander already knows the metaphorical views used, and therefore possesses source/target mappings underlying those views; however, the utterance uses aspects of the source domains(s) that are not mapped by the known mappings. The system is designed on the principle that there should by default be no attempt to create new mapping relationships to handle those aspects; rather, the system should try to do reasoning that links those aspects to source aspects that are already mapped. For example, consider again the sentence “*In the far reaches of her mind, Anne knew Kyle was having an affair.*” ATT-Meta handles this as follows, given knowledge of MIND AS PHYSICAL SPACE and IDEAS AS PHYSICAL OBJECTS, and most importantly the knowledge that (in)ability to

operate physically on an idea, in the source domain of IDEAS AS PHYSICAL OBJECTS, maps to (in)ability to operate mentally on the idea, in the target domain. Assume that ATT-Meta’s mappings do not map the far-reaches location within a space to anything, so that the utterance transcends the system’s mappings in this respect. ATT-Meta can reason, using (mainly) common-sense knowledge about physical spaces and objects, that Anne has only a very low ability to operate physically on the idea that Kyle was having an affair. (This is because the far reaches of a physical region are very distant from the main part of the region, and Anne, or rather her conscious self, is taken to be in that main part.) Then, using the known mapping, ATT-Meta can infer that Anne has only a very low ability to operate mentally on the idea. This example is treated in much more detail in Barnden & Lee (2001). A variety of other examples are also treated in that report and in other reports cited in it.

ATT-Meta’s long-term domain knowledge and its knowledge of metaphorical views is couched in terms of IF-THEN rules. A given metaphorical mapping relationship takes the form of a rule, such as (roughly)

*IF J is in reality an idea AND J is being viewed as a physical object AND person X is being viewed as being able to operate physically on J to at least degree D
THEN in reality X can operate mentally on J to degree at least D.*

ATT-Meta’s reasoning is entirely query-directed. So, for instance, in the Anne/Kyle example the reasoning steps mentioned arise from a backwards-going process of query construction, proceeding backwards through rules. In particular, a query about the degree of ability of Anne to operate mentally on the Kyle-affair idea in reality leads to the creation also of a query about the degree of ability of Anne to operate physically on the idea, under the metaphorical view. Thus, the system exhibits T-to-S transfer of queries.

The system’s metaphor-based reasoning is thoroughly integrated into a general-purpose framework for qualitatively uncertain reasoning. Reasoning in source-domain terms and in target-domain terms is generally uncertain. Rules and propositions are annotated with qualitative certainty levels, and there is a heuristic conflict-resolution mechanism that attempts to adjudicate between conflicting arguments. As a result of conflict-resolution, the certainty of one or more propositions is downgraded. Reasoning leaves behind a record of dependency links between propositions, so certainty downgrade of a proposition leads to downgrading also of propositions dependent on it. Now, for a given S→T mapping relationship there is often a converse mapping relationship, e.g. (to continue the above example),

IF J is in reality an idea AND J is being viewed as a physical object AND in reality person X can operate mentally on J to at least degree D

THEN X is being viewed as being able to operate physically on J to degree at least D.

Consequently, ATT-Meta performs T→S transfer of certainty downgrades when suitable target-domain propositions are downgraded. Because of the extensive within-source reasoning that ATT-Meta often performs, downgrade within the source domain can lead to other downgrades there by propagation along dependency links.

Because of the T→S mapping relationships, ATT-Meta can metaphorize literally-stated information, and such metaphorization steps are seamlessly mixed in with other reasoning steps. However, only limited experimentation on this has been done so far using ATT-Meta.

As for asymmetry of metaphor, it is instructive to look at the following situation that could hold in ATT-Meta. Recall the T→S rule mentioned above for IDEAS AS PHYSICAL OBJECTS. It can be reworded to have the overall form:

IF J is in reality an idea AND J is being viewed as a physical object AND ... THEN

Now consider S→T rules for a view of PHYSICAL OBJECTS AS IDEAS. These will have the overall form:

IF O is in reality a physical object AND O is being viewed as an idea AND ... THEN

The rule-forms differ crucially in their first two conditions. Thus, T→S rules for A-as-B will be act very differently from S→T rules for B-as-A, even when the source and target aspects mapped are the same or similar.

Conclusions

The main conclusions are that (a) target-to-source transfers of several distinctly different types are desirable in metaphorical discourse understanding; (b) in particular, metaphorization of within-target-domain utterances can be desirable; and (c) the directions of mapping relationships are sometimes intuitively the wrong way round in accounts of metaphor.

The ATT-Meta system is one that routinely allows target-to-source transfers of the sorts mentioned above. Blending theory (cited earlier) allows transfers into and out of the blended space, including transfers from the blended space back into a source domain. Also, the LISA model for analogy (Hummel & Holyoak, 1997) allows mapping to go from target to source, though the relationship to the considerations in the present paper is unclear. Thus, a small number of approaches are beginning to allow effects such as those in this paper. However, the topic appears to have seen little psychological experimentation or computational realization, and would be a fertile ground for future empirical investigation.

Acknowledgments

This research is supported by grant GR/M64208, Engineering and Physical Sciences Research Council.

References

- Barnden, J.A., Helmreich, S., Iverson, E., & Stein, G.C. (1996). Artificial intelligence and metaphors of mind: within-vehicle reasoning and its benefits. *Metaphor and Symbolic Activity*, 11 (2), 101–123.
- Barnden, 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), *Lecture Notes in Artificial Intelligence*, 1688. Springer.
- Barnden, J.A., & Lee, M.G. (2001). Understanding usages of conceptual metaphors: An approach and artificial intelligence system. Tech. Rep. CSRP-01-05, School of Computer Sci., Univ. of Birmingham, UK.
- Black, M. (1979). More about metaphor. In A. Ortony (Ed.), *Metaphor and Thought*. Cambridge, UK: Cambridge Univ. Press.
- Falkenhainer, B., Forbus, K.D., & Gentner, D. (1989). The Structure-Mapping Engine: algorithm and examples. *Artificial Intelligence*, 41 (1), 1–63.
- Gross, L. (1994). Facing up to the dreadful dangers of denial. *Cosmopolitan*, 216(3), USA ed.
- Hobbs, J.R. (1990). *Literature and cognition*. CSLI Lecture Notes, No. 21, Stanford University.
- Holyoak, K.J. & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13 (3), 295–355.
- Hummel, J., & Holyoak, K. (1997). Distributed representation of structure: A theory of analogical access and mapping. *Psychological Review*, 104 (3), 427–466.
- Johnson-Laird, P.N. (1983). *Mental models: towards a cognitive science of language, inference and consciousness*. Cambridge, MA: Harvard University Press.
- Lakoff, G. (1993). The contemporary theory of metaphor. In A. Ortony (Ed.), *Metaphor and Thought*, 2nd ed. Cambridge, UK: Cambridge Univ. Press.
- Markman, A.B. (1997). Constraints on analogical inference. *Cognitive Science*, 21 (4), 373–418.
- Ortony, A. (1979). The role of similarity in similes and metaphors. In A. Ortony (Ed.), *Metaphor and Thought*. Cambridge, UK: Cambridge Univ. Press.
- Turner, M., & Fauconnier, G. (1995). Conceptual integration and formal expression. *Metaphor and Symbolic Activity*, 10 (3), 183–204.
- Waggoner, J.E. (1990). Interaction theories of metaphor: psychological perspectives. *Metaphor and Symbolic Activity*, 5 (2), 91–108.
- Way, E.C. (1991). *Knowledge representation and metaphor*. Dordrecht: Kluwer.