

# Reasoning in Metaphor Understanding: The ATT-Meta Approach and System

John Barnden, Sheila Glasbey, Mark Lee, Alan Wallington  
School of Computer Science  
University of Birmingham  
Birmingham B15 2TT  
{J.A.Barnden, S.R.Glasbey, M.G.Lee, A.M.Wallington}@cs.bham.ac.uk

## Abstract

A detailed approach has been developed for core aspects of the task of understanding a broad class of metaphorical utterances. The utterances in question are those that depend on known metaphorical mappings but that nevertheless contain elements not mapped by those mappings. A reasoning system has been implemented that partially instantiates the theoretical approach. The system, called ATT-Meta, will be demonstrated. The paper briefly indicates how the system works, and outlines some specific aspects of the system, approach and the overall project.

## Introduction

The sentence *In the far reaches of her mind, Anne believed that Kyle was having an affair*<sup>1</sup> can be analyzed as depending on metaphorical views of MIND AS PHYSICAL SPACE and IDEAS AS PHYSICAL OBJECTS (see Barnden 2001a). These views are, plausibly, familiar to typical users of English. However, it is reasonable to assume that typical users do not already have a mapping into the mental domain of the physical notion of "far reaches". Our approach to metaphor is predicated on the notion that one should, when possible, avoid constructing source-to-target mappings for such elements of a metaphorical utterance that transcend the already known mappings in the metaphorical views underlying the utterance. Instead, we advocate the use of metaphor-pretence "cocoon" (reasoning spaces) where the utterance is taken as literally true. Within-

cocoon reasoning will attempt to connect these "map-transcending" elements to aspects of the source using a set of specified conversion rules. The *far-reaches* qualification in the sentence implies by source-domain reasoning that Anne could only to a very low degree physically manipulate the idea that Kyle was having an affair. Then, if we can appeal to a conversion rule, i.e. a known mapping of ability to physically manipulate an idea to ability to consciously process it, we can draw the reasonable conclusion that Anne only had a very low degree of ability to consciously process the idea. In our presentation, we shall demonstrate how the ATT-Meta system deals with this example.

Note that the rules of reasoning are given a qualitative certainty level, and that predicates can be graded, using a scale of qualitative degrees. For instance, someone can be represented as understanding a situation to a "medium" degree.

Our approach also makes heavy use of "view-neutral mapping adjuncts" (VNMA). These are general mapping principles (inspired by the work of Carbonell 1982) that apply, though only by default, no matter what metaphorical views are in play. For instance, the ability to do things and the degrees with which states of affairs hold are automatically mapped by VNMA. In many examples of metaphor, most of the real mapping work is done by VNMA.

Much of the approach has been implemented in the ATT-Meta system, which is an uncertain rule-based system operating by backchaining (see also Barnden 1998, Barnden 2001, Lee &

---

<sup>1</sup> Slightly adapted from real discourse.

Barnden, 2001a). ATT-Meta performs reasoning, but does not yet interface directly to natural language. Instead, hand-constructed logical forms couching the source-domain meaning of metaphorical sentences are passed to it. In the above example, the source-domain meaning is that Anne's believing was literally physically located in the physical far reaches of her mind.

The following sections summarize various abilities of the system, principles of the approach, and aspects of ongoing theoretical work aimed at further extensions to the system. A major item of current implementational work is a fuller realization of VNMA's.

## 1 Uncertainty

Although reasoning conflict and uncertainty are intricately involved in metaphor, very few approaches attempt to grapple with the issues. Propositions and reasoning within both the target and source domains, being largely of common-sense varieties, are typically uncertain. It can be uncertain what metaphorical views are involved; information transferred from the source domain can conflict with target-domain information; and transfers can even conflict with each other. The ATT-Meta system handles all these types of uncertainty and conflict. Its uncertainty handling is based on fairly crude qualitative uncertainty annotations on rules and propositions, but there is a sophisticated conflict-resolution mechanism.

The uncertainty-handling and conflict-resolution are almost entirely orthogonal to the provisions for metaphor. This leads to clean design and helps to address long-standing issues about metaphor. One such issue is the conflict between information transferred from the source domain and the target information. ATT-Meta allows either side to win, depending on standard specificity principles. This goes against a naive assumption in most of the literature that target information should automatically override transfers. But, this is only convincing when the target information is certain. Indeed, we claim that metaphor is often used precisely to describe an exception to a target-domain default.

## 2 Mixed Metaphors

Issues such as reasoning about uncertainty are particularly important in the processing of mixed metaphors. Mixed metaphors need not feature obvious cases of conflict but can include graceful combinations of metaphors, such as the following sentence to be examined below: *One part of John hotly resented the verdict*. This combines a view of John as made up of sub-agents and a view of agents' emotional states as things that can have temperature. It is possible to distinguish two types of mixed metaphor: parallel mixes and serial mixes. In a parallel mixed metaphor, the target (A) is seen partly through an A-as-B metaphor and partly through another metaphor, A-as-B'. B and B' are in general different domains, but may overlap. Also, different aspects of A may be involved in the two metaphors. In a serial mixed metaphor (commonly called a chained metaphor), the target (A) is seen as a source (B), which is in turn then seen as a different source (C).

Previous work on the understanding of metaphor has assumed that mixing is a relatively rare phenomenon that can be handled once a more theory of simple metaphor is developed. We argue that this assumption is detrimental to progress since mixed metaphors rely on the same conceptual knowledge as simple metaphors and can, therefore, provide valuable insight into the processes and representations underlying metaphorical reasoning. Moreover, we claim that the reasoning processes and data structures involved in understanding mixed metaphors are identical to those used in understanding simple metaphors. Therefore, any current theory of metaphor should (at least in principle) be extensible to deal with mixing. To this end, ATT-Meta handles mixed metaphor in a manner consistent with the way it handles simple metaphors. The two types of metaphor are processed in subtly different ways. Parallel mixed metaphors create separate pretence-cocoons that are mapped in parallel to the target domain where their respective contributions are understood. Serial mixed metaphors create nested pretence cocoons where the metaphorical view of B as C is nested within a pretence cocoon with the view of A as B.

### 3 Reverse Transfers in Metaphor

The use of metaphor involves a flow of effects of some kind from the source domain to the target domain, where effects can include insights into the target, hypotheses about the target, or the highlighting of parts of the target. However, although the overall effect flow is always from source to target, in many cases, this does not preclude a reverse flow where a literal proposition, command, or question is mapped onto an equivalent within the current metaphorical domain.

The ATT-Meta system allows conversion rules to map from propositions in the source domain to propositions in the target domain and also in the opposite direction. So a source domain proposition such as "Socrates was the midwife for an idea" might be mapped onto the target domain proposition "Socrates helped in the production of the idea". However, the rules would equally allow the proposition "Socrates helped in the production of the idea" to be mapped to the source domain proposition "Socrates was the midwife for an idea". We argue that there are at least three reasons why ATT-Meta should have this ability:

(1) Given that metaphors are ultimately used to have an effect on the target domain, the use of a metaphorical utterance can be seen as answering, in some sense, a target domain query. This sets up a choice between taking the metaphorical utterance and applying all conversion rules to it in the hope that one of the resulting propositions might provide a suitable answer, or taking the question and converting it into a question in terms of the current metaphor. We argue that the latter is often more efficient.

(2) Certain source domain propositions would allow ATT-Meta to draw a tentative conclusion, which would, were it more strongly supported, provide an argument via a chain of reasoning for some other, target level, proposition or query. A target-level statement might give the added support, but for this to be the case it would first need to be converted into its source-level equivalent.

(3) The combination of source and target domain information within a discourse that only intermittently maintains a metaphorical view of the target domain may best be done in the source domain after the target domain information has been "metaphorized". This would be especially so if the source domain was information-rich compared to the target domain, so allowing much more reasoning to be carried out than would be possible in the target domain.

### 4 Non-Declarative Metaphor

Almost all examples of metaphorical language discussed in the literature are of declarative utterances rather than questions, commands, ejaculations, etc. However, these other forms of utterance can obviously occur. For instance, just as one can state "John is a steamroller" one can ask "Is John a steamroller?" Just as one can state "The champion knocked the cream-puff out" one can issue the command "Knock that cream-puff out!" The observation that questions, in particular, can be metaphorical, plays a significant role in our theoretical approach. This is because their processing is contiguous with that of implicit queries generated within the metaphorical pretence cocoon (see Introduction) during ATT-Meta's goal-directed reasoning. However, the theoretical significance of non-declarative metaphorical utterances is even greater, because such utterances call into question accounts of metaphor that assume the task of understanding is to work out what claim about the target domain the metaphorical utterance is making.

Compiling such examples is an additional goal of our corpus work (see section 6).

### 5 Time and change

Work is ongoing which addresses the temporal, aspectual and causal facets of metaphor. A survey of metaphors in the ATT-Meta Databank reveals, unsurprisingly, that the metaphorical expressions there involve a wide range of tense and aspectual constructions in English, including past, present and future tenses, simple and progressive aspects, and the full set of aspectual classes. A wide variety of temporal adverbials is also present. A key topic under investigation is

the mapping of temporal and aspectual information between source and target domains. For example, if an event is telic in the source domain, to what extent does that telicity carry over to the target domain? Preliminary investigations confirm the expectation that such aspectual information is preserved in the majority of cases. Exceptions exist, however, and these merit further study.

The mapping of temporal duration between domains is also being investigated. In some cases, a mapping appears to exist whereby an event of long duration relative to the source domain maps to an event with long duration relative to the target domain. This can be captured by an appropriate VNMA, which maps relative durations between domains. The logic of ATT-Meta is episode-based, which means that it is relatively straightforward to express this kind of constraint and employ it in reasoning. Currently underway is a detailed examination of metaphorical expressions involving both explicit and implicit temporal durations. This will result in a set of VNMA's covering a wide range of tense/aspect/temporal-adverbial constructions.

A second strand of the work on time involves a detailed study of the metaphors used to describe times, states and events, including spatial metaphor for time (Lakoff 1994).

## **6 Corpus Studies of Metaphor**

As an adjunct to the development of the ATT-Meta approach and system, we have been conducting corpus studies of metaphor, mainly using the British National Corpus but also using the Bank of English and, to a limited extent, web search engines. We have used both hand-annotation of small numbers of documents from the BNC and automated search for particular types of metaphorical phraseology (mainly relatively fixed metaphorical phrases concerning mental states) over the whole of corpora.

Current objectives are (a) to develop large databanks of examples of various types of metaphorical utterance, for the benefit of metaphor researchers in general, (b) to demonstrate more extensively and objectively

the importance in discourse of "map-transcending" metaphorical utterances (see Introduction), (c) relatedly, to reveal the degree to which relatively conventional metaphor phraseology can be varied in real discourse (cf. Moon 1998), and (d) to uncover (in small numbers of documents) the degree to which metaphorical utterances relate to context: how much their understanding depends on context and how much the understanding of the context depends in turn on them. We are interested in (d) because in the ATT-Meta approach the process of metaphorical understanding is partially guided by discourse goals set up by context. This feature goes a long way to side-stepping problems of apparent indeterminacy of meaning of metaphorical utterances when taken in isolation.

We also have the methodological objectives of developing a good annotation regime for metaphor and better-automated search techniques for metaphor. As part of the latter, we plan to investigate the usefulness of a large set of morphological, lexical, syntactic and phraseological clues to the presence of metaphor, inspired by the clues discussed by Goatly (1997). These clues are only present in a minority of metaphorical utterances but could nevertheless form a useful weapon in the automated search armoury.

Metaphor detection techniques developed for corpus study should also help with developing a means for an understanding system to notice the presence of metaphor. Such noticing is not currently performed by ATT-Meta but is an important topic for future research.

## **Conclusion**

The ATT-Meta project is making headway in showing how metaphorical utterances can be computationally processed. It is based on a distinctive set of principles as to how to understand metaphor, some of which are original and some related to those of previous researchers. In particular, it seeks to avoid expensive computation of new analogical mappings between domains as a regular part of metaphorical understanding. This is inspired

partly by the observation that genuinely novel pairings of domains are relatively rare in real discourse. What are more common are novel extensions of familiar metaphorical views, and novel mixes of views. This is true even in poetry (Lakoff and Turner 1989). The project is also seeking to take full account of the important role that uncertainty, gradedness and dynamism of situations plays in metaphor.

The approach and system have been evaluated in a number of ways. We have applied the implemented system or the theoretical approach to (simplified versions of) selected real-discourse examples from an existing databank (<http://www.cs.bham.ac.uk/~jab/ATT-Meta/Databank>): see Barnden (2001a), Barnden & Lee (2001a) and Barnden & Lee (2001b). We have applied the implemented system to examples of all the metaphors of mental states listed in the Master Metaphor List (Lakoff 1994, Lee & Barnden 2001a). The examples here were found by search over the Bank of English. Finally, we have applied the theoretical approach to various real-discourse examples included in Goatly (1997): see Barnden (2001b).

## Acknowledgements

This research is supported by grant GR/M64208 from the Engineering and Physical Sciences Research Council of the UK.

## References

- Barnden, J. A. (1998a) *Combining uncertain belief reasoning and uncertain metaphor-based reasoning*. In "Procs. Twentieth Annual Meeting of the Cognitive Science Society", Lawrence Erlbaum Associates, Hillsdale, N.J, pp.114-119.
- Barnden, J.A. (2001a) *Application of the ATT-Meta Metaphor-Understanding Approach to Various Examples in the ATT-Meta Project Databank*. Technical Report CSRP-01-02, School of Computer Science, The University of Birmingham, U.K.
- Barnden, J.A. (2001b) *Application of the ATT-Meta metaphor-understanding approach to selected examples from Goatly*. Technical Report CSRP-01-01, School of Computer Science, The University of Birmingham, U.K.
- Barnden, J.A (2001c) *Uncertainty and conflict handling in the ATT-Meta context-based system for metaphorical reasoning*. In "Lecture Notes in Artificial Intelligence", J.G.Carbonell and J.Siekmann, eds, Springer Verlag, Berlin.
- Barnden, J.A. and Lee, M.G. (1999) *An implemented context system that combines belief reasoning, metaphor-based reasoning and uncertainty handling*. In "Modelling and Using Context", P. Bouquet, P. Brezillon, L. Serafini ed., Lecture Notes in Artificial Intelligence, No. 1688, Springer-Verlag, pp. 28-41.
- Barnden, J.A. and Lee, M.G. (2001a) *Understanding open-ended usages of familiar conceptual metaphors: An approach and artificial intelligence system*. Technical Report CSRP-01-05, School of Computer Science, The University of Birmingham, U.K.
- Barnden, J.A. and Lee, M.G. (2001b) *Application of the ATT-Meta metaphor-understanding system to an example of the metaphorical view of MIND PARTS AS PERSONS*. Technical Report CSRP-01-09, School of Computer Science, The University of Birmingham, U.K.
- Carbonell, J.G. (1982) *Metaphor: an inescapable phenomenon in natural-language comprehension*. In "Strategies for Natural Language Processing", W. Lehnert & M. Ringle ed., Lawrence Erlbaum, Hillsdale, N.J, pp.415-434.
- Goatly, A. (1997) *The language of metaphors*. Routledge, London and New York.
- Lakoff, G. (1994) *The Master Metaphor List*. <http://cogsci.berkeley.edu/>, University of California, Berkeley.
- Lakoff, G. and Turner, M. (1989) *More than cool reason: a field guide to poetic metaphor*. University of Chicago Press, Chicago.
- Lee, M. G. and Barnden, J. A. (2001a) *Reasoning about mixed metaphors within an implemented AI system*. *Metaphor and Symbol* 16, 1& 2, pp. 29-42.
- Lee, M. G. and Barnden, J. A. (2001b) *Mental Metaphors from the Master Metaphor List: Empirical Examples and the Application of the ATT-Meta System* Technical Report CSRP-01-03, School of Computer Science, The University of Birmingham, U.K.
- Moon, R. (1998) *Fixed idioms and expressions in English*. Clarendon Press, Oxford, U.K.