

# Combining Uncertain Belief Reasoning and Uncertain Metaphor-Based Reasoning

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

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

## Abstract

An implemented AI reasoning system called ATT-Meta is sketched. It addresses not only AI issues but also ones that are salient in psychology, philosophy, cognitive linguistics, discourse pragmatics and other disciplines. These issues include the Simulation-Theory/Theory-Theory debate and Fauconnier and Turner's notion of conceptual blending. The system performs metaphor-based reasoning and reasoning about mental states of agents; in particular, it performs metaphor-based reasoning about mental states. Although it relies on built-in knowledge of specific conceptual metaphors, it is flexible in allowing novel discourse manifestations of those metaphors. The metaphorical reasoning and mental-state reasoning facilities are fully integrated into a general framework for uncertain reasoning. A special result of the overall approach is that it enables a unified handling of certain apparently separate discourse phenomena: chained metaphor, personification metaphor, and reports of agents' own metaphorical thoughts.

## Introduction

There are two main strands in this paper: *belief reasoning* (reasoning about the beliefs and attendant reasoning of other agents, including the case when those beliefs etc. are themselves about beliefs etc.); and *metaphor-based reasoning*. These topics are normally studied separately, but of course metaphorical utterances can be about the mental states and processes of agents. The implemented system (ATT-Meta) on which this paper is centered is mainly geared towards such utterances, but it can also do non-metaphor-based belief reasoning and metaphor-based reasoning about non-mental topics. (The name "ATT-Meta" comes from "[propositional] Attitudes" and "Metaphor.")

A third research strand is uncertainty-handling. This is important both for useful, commonsensical belief reasoning and useful, commonsense reasoning based on metaphor. Yet, it is relatively uncommon for schemes for belief reasoning, especially implemented ones, to involve an extensive treatment of uncertainty (though see: Asher & Lascarides, 1994; Chalupsky, 1996; Cravo & Martins, 1993; Dragoni & Puliti, 1994; and Parsons, Sierra & Jennings, 1998). Similarly, although metaphor researchers often make informal reference to uncertainties, Hobbs (1990) is one of very few to devote extensive technical attention to uncertainty. (Indeed, he is one of the very other researchers to provide a detailed computational framework for metaphor-based reasoning, as opposed to a framework for deriving meanings of metaphorical utterances.) One goal of the ATT-Meta research has been to have the system perform uncertain belief reasoning and uncertain metaphor-based reasoning in a systematic way, doing more

justice to uncertainty than has heretofore been seen. ATT-Meta's treatment of uncertainty is still only a first approximation to what is needed, but the work draws attention to crucial but unrecognized complications in belief reasoning and metaphor-based reasoning.

For the above and other reasons, this AI research is of broad interdisciplinary relevance. It connects to psychology, philosophy, cognitive linguistics, discourse pragmatics and other areas. The research brings in the following particular issues, among others, in those areas: the Simulation-Theory/Theory-Theory debate (see, e.g., Davies & Stone, 1995, Carruthers & Smith, 1996), the unparaphrasability of much metaphor usage, the role of literal meaning in metaphorical processing, the overriding of metaphor-based inferences by tenor information or vice versa, the role of metaphor in thought as well as language, and conceptual blending (Turner & Fauconnier, 1995).

An early version of ATT-Meta was reported in a previous annual conference of the Cognitive Science Society (Barnden *et al.*, 1994). The current system includes major algorithmic and conceptual advances over that version. The mechanisms for conflict resolution have been refined and greatly extended, so as to work properly across multiple layers of belief and metaphor. Certain interactions between belief reasoning and metaphor-based reasoning have been streamlined and made more general. But in any case the present paper focuses on different aspects of the research from those stressed in the earlier paper.

The remaining sections of the paper are as follows: a section on the main type of metaphorical utterance considered in the research; a section very briefly sketching ATT-Meta's basic reasoning facilities and uncertainty-handling, irrespective of metaphor; a section describing ATT-Meta's metaphorical reasoning; a section on various types of uncertainty handled in the metaphorical reasoning; a section sketching the facilities for reasoning about agents' beliefs and reasoning, irrespective of metaphor; a section on how belief reasoning and metaphor-based reasoning can interact in ATT-Meta; a section on the similarities and differences between ATT-Meta's belief reasoning and its metaphor-based reasoning; a section on connections to some topical issues; and a brief conclusion section.

As a major function of the paper is to stress interdisciplinary connections, it will not go into fine technical detail. Further detail of the work is included in the Barnden (*et al.*) items in the References section.

ATT-Meta is merely a reasoning system, and does not deal

with natural language input directly. Rather, a user supplies hand-coded logic formulae that are intended to couch the literal meaning of small discourse chunks (two or three sentences).

### Metaphor in ATT-Meta

A metaphorical utterance is one that *manifests* (instantiates) a metaphor, where a *metaphor* is a conceptual view of one topic as another. Here I broadly follow Lakoff (e.g., Lakoff, 1993). An example of a metaphor is the view of the mind as a three-dimensional physical region (MIND AS PHYSICAL SPACE). A metaphor is the *view* itself, as opposed to some piece of natural language that manifests the view. Such a manifestation might be "John believed in the recesses of his mind that ...," in the case of MIND AS PHYSICAL SPACE. In a manifestation, the topic actually being discussed (John's mind, in the example) is the *tenor*, and the topic that it is metaphorically cast as (physical space, in the example) is the *vehicle*.

The ATT-Meta system does not currently deal with novel *metaphors* — rather, it has pre-given knowledge of a specific set of metaphors, including MIND AS PHYSICAL SPACE. But it is specifically designed to handle novel *manifestations* of those metaphors. Its knowledge of a metaphor consists mostly of a relatively small set of very general "conversion rules" that map between the vehicle and potential tenors. The degree of novelty that the system can handle in a manifestation of a metaphor is limited only by the amount of knowledge it has about the vehicle and by the generality of the conversion rules. Note also Lakoff & Turner's (1989) persuasive claims that even in poetry metaphorical utterances are mostly manifestations of familiar, well-known metaphors, albeit the manifestations are highly novel and metaphors can be mixed in novel ways.

The metaphor research underlying ATT-Meta has concentrated on metaphors for mental states, such as MIND AS PHYSICAL SPACE, although the principles and algorithms implemented are not restricted to such metaphors. Mundane discourses, such as ordinary conversations and newspaper articles, often use metaphor in talking about mental states/processes of agents. Indeed, as with many abstract topics, as soon as anything at all subtle or complex needs to be said, metaphor is practically essential. There are many mental-state metaphors apart from MIND AS PHYSICAL SPACE. Some are as follows: IDEAS AS PHYSICAL OBJECTS, under which ideas are cast as physical objects that have locations and can move about, as in "He pushed these ideas to one side;" COGNITION AS VISION, as when understanding, realization, knowledge, etc. is cast as vision, as in "His view of the problem was blurred;" IDEAS AS INTERNAL UTTERANCES, which is manifested when a person's thoughts are described as internal speech or writing (internal speech is not *literally* speech), as in "He said to himself that he ought to stay at home and work;" and MIND PARTS AS PERSONS, under which a person's mind is cast as containing several sub-agents with their own thoughts, emotions, etc., as in "Part of him was convinced that he should go to the party." Many real-discourse examples of mental-state metaphor can be found in a databank at <http://www.cs.bham.ac.uk/~jab/ATT-Meta/Databank>.

### ATT-Meta's Basic Reasoning

ATT-Meta is a rule-based reasoning system that manipulates hypotheses (facts, conclusions or goals), represented as expressions in a situation-based/episode-based first-order logic somewhat akin to that of Hobbs (1990). At any time, any particular hypothesis H is tagged with a certainty level, one of certain, presumed, suggested, possible or certainly-not. The last just means that the negation of H is certain. Possible just means that the negation of H is not certain but no evidence has yet been found for H itself. Presumed means that H is a *default*: i.e., it is taken as a working assumption, pending further evidence. Suggested means that there is evidence for the hypothesis, but it is not (yet) strong enough to enable H to be a working assumption.

ATT-Meta applies its rules in a backchaining style. It is given a reasoning goal, and uses rules to generate supporting goals. Goals can of course also be satisfied by provided facts. When a rule application supports a hypothesis, it supplies a level of certainty to the hypothesis, calculated as the minimum of the rule's own certainty level and the levels picked up from the hypotheses satisfying the rule's condition part. When several rules support a hypothesis, the maximum of their certainty contributions is taken.

When both a hypothesis H and its negation  $\neg H$  are supported to level at least *presumed*, *conflict-resolution* takes place. The most interesting case is when both hypotheses are supported to level *presumed*. The system attempts to see whether one hypothesis has *more specific* evidence than the other, so that it can downgrade the certainty level of the other hypothesis. Specificity comparison is a commonly used heuristic for conflict-resolution in AI, although serious problems remain in coming up with adequate and practical heuristics. ATT-Meta's specificity comparison is closely related to other schemes in the literature. Under certain conditions, one way for a hypothesis to be more specifically supported than its negation is for it to be supported (directly or indirectly) by a proper superset of the facts supporting the negation. Inter-derivability relationships between hypotheses appearing in the support networks are also used in specificity comparison.

If a hypothesis is more specifically supported than its negation, it stays *presumed* and the negation is downgraded to *suggested*. If neither hypothesis wins, both are downgraded to *suggested*. The scheme can deal with any amount of iterative defeat: for example, if magic penguins are special penguins that can indeed fly, but ill magic penguins once again cannot fly, then the system will resolve the conflicts correctly for magic penguins in general and for ill magic penguins.

This paper will not display ATT-Meta's formal representations and formal rule formats (which are in turn represented as Quintus Prolog expressions), and will use English glosses instead. These glosses may use the past tense to match the tense of English example sentences, but this is just for readability, and ATT-Meta currently has no treatment of time.

### Metaphor-Based Reasoning

Notoriously, metaphorical utterances can be difficult if not impossible to paraphrase in non-metaphorical terms. Similarly, it can be difficult if not impossible to give them internal

meaning representations that are not themselves metaphorical. Consider, for instance, “One part of John was insisting that Sally was right.” This manifests the metaphor of MIND PARTS AS PERSONS, where furthermore the mentioned “part” engages in natural language communication. We simply do not know enough about how the mind works to give, in non-metaphorical terms, a useful and reasonably full account of what was going on in John’s mind according to the sentence. What useful non-metaphorical account can be given of some “part” of John “insisting” something? Rather, the utterance connotes things such as the following:

### Connotation

*John had reasons both to believe that Sally was right and to believe the opposite.*

This particular connotation arises because *a person generally insists something only when someone else has stated the opposite* (although there are other possible scenarios). So, the sentence suggests that

*some other “part” of John stated, and therefore probably believed, that Sally was NOT right.*

Then, because of the thoughts of the two sub-agents within John (the two parts), we can infer the connotation displayed above. This is assuming that the system, as part of its knowledge of the MIND PARTS AS PERSONS metaphor, knows that

((K)) *if a “part” of someone X believes something P, then X has reasons to believe P.*

Some investigators may wish to call the above connotation the underlying *metaphorical meaning* of the utterance, or at least to claim it to be part of that meaning. In contrast, I have avoided the difficult task of defining the notion of “metaphorical meaning,” and have concentrated instead on the broader question of algorithms for making commonsense inferences, by whatever route, from metaphorical utterances. This is liberating, because the question of which sector of the space of possible inferences should be called the metaphorical meaning is merely a terminological issue. I assume instead that a metaphor-based reasoning system should, in many cases at least, construct a *literal* meaning of the metaphorical utterance in question, and should make inferences from it. (The literal meaning for the above example sentence is that John literally had a part that literally insisted that Sally was right.) Some of those inferences will themselves be couched in metaphorical terms — even though they are internally represented, rather than represented in natural language — and some will be in non-metaphorical terms; but the latter can be at an arbitrary inferential distance from the utterance. Because of the attitude adopted towards metaphorical meaning, we can say that ATT-Meta is “semantically agnostic” as regards metaphor. The approach is akin to but less extreme than that of Davidson (1979), which can be regarded as semantically “atheist.”

ATT-Meta’s approach to deriving connotations such as the one above is *literal pretence*. A literal-meaning representation for the metaphorical input utterance is constructed. The system then *pretends* that this representation, however ridiculous in reality, is true. Within the context of this pretence, the

system can do any reasoning that arises from its knowledge of the vehicles of the metaphors involved. In our example, it can use knowledge about interaction within groups of real people, and knowledge about communicative acts such as insistence. As a result of this knowledge, the system can infer, within the pretence, that the explicitly mentioned part of John believed (as well as insisted) that Sally was right, and some other, unmentioned, part of John believed (as well as stated) that Sally was not right. These conclusions are examples of internal metaphorical hypotheses.

The key point is that this reasoning from the *literal* meaning of the utterance, conducted within the pretence, link up with the knowledge displayed as (K), which maps from the vehicle to the tenor of the metaphor. That knowledge is itself of a very fundamental, general nature, and does not, for instance, rely on the notion of insistence or any other sort of communicative act. Any line of within-pretence inference that linked up with that knowledge could lead to conclusions that John had reasons to believe certain things. This is the way in which ATT-Meta can deal with novel manifestations of metaphors. There is no need at all for it to have any knowledge of how “insistence” by a “part” of a person maps to some non-metaphorically describable property of the person. Thus, the system needs no prior exposure to, or specific knowledge of how to deal with, metaphorical utterances involving “insistence” by parts of people.

To implement the sketched approach on the “insisting” example, ATT-Meta proceeds as follows in dealing with the logical input corresponding to a metaphorical utterance. This input includes an encoding — (L) below — of the literal meaning of the utterance. The system constructs a computational environment called a *metaphorical pretence cocoon*. The following shows hypotheses that are placed inside and outside the cocoon:

---

### Inside the Cocoon

((L)) A part PJ of John insisted that Sally was right.  
 ((PJ)) PJ is a person.

---

### Outside the Cocoon

((SL)) I (the system) am pretending that (L) holds.  
 ((SPJ)) I (the system) am pretending that PJ is a person.

---

As usual, the system is given a reasoning goal, such as

((G1)) John believed Sally was right.

Assume the system has

### Rule R:

IF X has reasons to believe P  
 THEN [presumed] X believes P.

The “presumed” is the rule’s certainty qualifier, and has the effect of limiting any conclusion of the rule to be at best presumed. In application to goal (G1), the system gets the subgoal

((G2)) John had reasons to believe that Sally was right.

Now, knowledge item (K) appears in the system as the following “conversion” rule, converting between metaphorical and non-metaphorical terms:

### Conversion Rule KCR

```
IF I (the system) am pretending that
  part Y of agent X is a person
AND I am pretending that Y believes Q
THEN [presumed] X has reasons to believe
  Q.
```

In application to (G2), the rule leads to the creation of the subgoal

```
((G3)) I (the system) am pretending that
  PJ believed Sally was right.
```

All the goals so far mentioned are outside the metaphorical pretence cocoon, but (G3) is automatically accompanied by the subgoal

```
((G4) PJ believed that Sally was right
```

*within* the cocoon. This hypothesis can then be inferred (as a default) from the hypothesis that PJ stated that Sally was right, which itself can be inferred (as a default) from the existing within-cocoon fact (L). Notice carefully that these last two steps are *entirely within the cocoon and merely use commonsense knowledge about real-life communication*.

As well as the original goal (G1) the system also looks at the negation of (G1)-and therefore also at the hypothesis that John believed that Sally was *not* right, and hence, because of Rule R, at the hypothesis that John had reasons to believe that Sally was not right. This subgoal gets support in a rather similar way to the above process, but it involves richer reasoning within the cocoon.

### Uncertainty in Metaphor

A hypothesis like “I (the system) am pretending that P” is called a *pretence hypothesis*. In our example, such a formula arises outside the cocoon mentioned above. When it does, a copy of “P” is placed inside the cocoon. Conversely, every hypothesis P that arises within the cocoon is reflected outside by the corresponding pretence hypothesis. The hypotheses within the cocoon are noted as being within the cocoon by being tagged with the system’s name for the cocoon. Such tags are passed around by reasoning rules, so that rule applications on hypotheses within the cocoon lead only to within-cocoon hypotheses. But the tags do not otherwise affect rule application. Thus, application of a rule within a cocoon is virtually identical to application outside the cocoon. (And, currently, all rules available for the system’s reasoning outside cocoons can also be used within cocoons.) In particular, uncertainty is handled within the cocoon just as it is outside.

ATT-Meta includes the following three types of uncertainty handling in its metaphor-based reasoning.

(UM1) Given an utterance, it is often not certain what particular metaphors or variants of them are manifested. Correspondingly, ATT-Meta may merely have *presumed*, for instance, as a tentative level of certainty for a pretence premise

like (SPJ) above. This hypothesis is then potentially subject to defeat.

(UM2) Conversion rules like KCR 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 specificity of the evidence for and against the conclusion. Thus, whether a piece of metaphorical reasoning overrides or or is overridden by other lines of reasoning about the tenor is matter of the peculiarities of the case at hand. However, many researchers (e.g., Lakoff 1993) assume that, in cases of conflict, tenor information should override metaphor-based inferences, and thus do not fully address the potential uncertainty of tenor information. It must be realized that, just as with literal utterances, a metaphorical utterance can express an exception to some situation that would *normally* apply in the tenor domain. To say “The company nursed its competitor back to health” contradicts default knowledge that companies do not normally help their competitors, and should override that knowledge.

(UM3) Knowledge about the vehicle of the metaphor is itself generally uncertain. Correspondingly, in ATT-Meta the hypotheses and reasoning within the cocoon are usually uncertain. For instance, it is not certain that someone believes something just because they state it. (A default step from stating to believing was used in the “insisting” example.)

Because there is uncertain reasoning both within and outside the cocoon, special complications arise for conflict resolution. A particular complication is that the pretence cocoon is taken to contain as a fact any fact sitting outside. This importation of facts is needed because arbitrary information about, say, physical objects may be needed in a pretence cocoon used for a metaphor like MIND AS PHYSICAL SPACE. Within the cocoon, the imported facts may support something that conflicts with conclusions drawn from the special metaphorical facts inserted into the cocoon at the start (e.g., the fact (PJ) that part PJ of John is a person). However, the system adopts the heuristic that such metaphorical facts supply added specificity. Therefore, ATT-Meta proceeds as follows: within a metaphorical pretence cocoon, specificity-comparison is first attempted in a mode where all reasoning lines partially dependent on imported facts are thrown away. Only if this does not yield a winner are those lines restored, and specificity reassessed.

### Agents’ Beliefs and Reasoning

ATT-Meta can reason non-metaphorically about the beliefs and reasoning acts of agents, to any depth of nesting of agents. Although ATT-Meta can reason about beliefs in an ordinary rule-based way, its main tool is *simulative reasoning* (see, e.g.: Haas, 1986; Dinsmore, 1991; Chalupsky, 1996). In attempting to show that *agent X believes P (to some specific level of certainty)* from the fact that *X believes Q (to some level)*, the system puts P as a goal and Q as a fact in a “simulation cocoon” for X, which is a special environment which is meant to reflect X’s own alleged reasoning processes. Reasoning from Q to P in the cocoon is alleged (by default) to be reasoning by X.

Belief reasoning in ATT-Meta can involve uncertainty in the following three important senses:

(UB1) Information gained from mental state reports in discourse can be uncertain, because of hedges in sentences or speaker unreliability. Correspondingly, in ATT-Meta a hypothesis of the form *X believes that H (to some level)* can itself be uncertain.

(UB2) Even if belief facts like *X believes ...* in (UB1) were certain, further conclusions drawn from them about X's beliefs must generally be uncertain. If those conclusions are reached by ordinary rule application, the rules can be uncertain. Also, simulation of X never supports a hypothesis of form *X believes that G to some level* to a level higher than presumed. This is to allow for the point that the agent doing the simulating cannot be certain that X does the alleged steps, and in any case X may perform unknown steps that provide an argument against G. In addition, simulative and non-simulative belief reasoning about an agents' beliefs can conflict. Conflict resolution might resolve the conflict either way, depending on circumstances.

(UB3) The reasoning within the simulation is itself generally uncertain. It can involve uncertain rule-based reasoning, and it can involve simulation of further agents.

### Belief/Metaphor Interactions

Belief reasoning and metaphor-based reasoning can interact in a variety of ways. One simple way is illustrated by the sentence "John thought, 'Bill is a fool'," which is a manifestation of IDEAS AS INTERNAL UTTERANCES (see Barnden *et al.*, 1996 and Barnden, in press, for discussion). A simple inference from this (via a conversion rule application to the literal meaning) is that John believed that Bill was a fool.

A more indirect type of connection is the main point of Barnden *et al.* (1994). In the sentence "These two ideas were far apart in John's mind," which manifests MIND AS PHYSICAL SPACE, the connotation is that John did not draw otherwise expectable inferences from the two ideas. Whereas ATT-Meta would normally simulate John as making the inferences, there is a mechanism whereby the far-apartness blocks the simulation from having its normal effect. In brief, each reasoning step in a simulation is accompanied, outside the simulation, by a hypothesis to the effect that the agent does the step. These special hypotheses can be reasoned about just as any other hypothesis can, and in particular they can be defeated in the face of special evidence like our example metaphorical sentence.

Another type of metaphor/belief interaction occurs in *personification* metaphor, as in "My car thinks it's Sunday" uttered to explain why the car won't start. From the fact that the car thinks it's Sunday, we could infer that the car thinks it needn't "wake up" until some relatively late time. This alleged reasoning by the car would occur within a simulation cocoon for the car, embedded within a metaphorical pretence cocoon for the pretence that the car is a person. Indeed, the example sentence "One part of John was insisting that Sally was right" used above is also a manifestation of personification metaphor.

Conversely, metaphorical pretence can be embedded within simulative reasoning about beliefs. An agent X that is mentioned in discourse may be portrayed as thinking and reasoning metaphorically about something, as in (one interpretation of) the sentence "Romeo thinks that Juliet is the sun."

This phenomenon can be handled by embedding a metaphorical pretence cocoon for pretending that Juliet is the sun within a simulation cocoon for simulating Romeo.

### Belief/Metaphor Reasoning Similarity

ATT-Meta's metaphorical pretence reasoning and simulative belief reasoning are almost identical. Note in particular that uncertainty types (UB1) to (UB3) are directly analogous to uncertainty types (UM1) to (UM3), respectively. However, there are some major points of difference algorithmically:

(a) Simulative belief reasoning does not *need* an analogue of conversion rules adopted for metaphors. One can imagine, however, directly analogous rules of the form *IF an agent of such-and-such a type believes P AND P is about so-and-so THEN [presumed] P* or of the converse form *IF P AND P is about so-and-so and agent X is of such-and-such a type THEN [presumed] X believes P*. The former would sanction reliance on the opinions of trusted others, and the latter would sanction default ascription of beliefs to suitable agents (cf. Ballim & Wilks, 1991).

(b) The above-mentioned automatic importation of facts into a metaphorical pretence from outside has no analogue in simulative belief reasoning.

We have seen that belief reasoning and metaphorical pretence can be nested within simulative belief reasoning, and belief reasoning can be nested within metaphorical pretence. Also, we have the fourth possibility, namely that metaphorical pretence can be nested within metaphorical pretence. This handles *chained metaphor* (A-AS-B mixed with B-AS-C). Consider the sentence "The thought hung over him like an angry cloud" (adapted from a real-text example). The thought is metaphorically cast as a cloud, and the cloud is in turn metaphorically cast as an animate being (because only animate beings can literally be angry). This can be handled by having a metaphorical cocoon for the second of those two metaphorical steps nested within a cocoon for the first. That is, within the pretence that the thought is a cloud there is a further pretence that the cloud is a person.

In principle, the four types of nesting can be done to arbitrary depth. This fact, in conjunction with uncertainty types (UM1-3) and (UB1-3), requires conflict resolution to be handled correctly within different cocoons at any depth of nesting, and across the boundaries of cocoons. The latter is needed because, outside a particular cocoon (for pretence or ordinary simulation), a hypothesis might be supported and attacked by a combination of argumentation based on that cocoon and other argumentation. The resultant complications in conflict resolution appear not to have been addressed elsewhere. The mechanism used in ATT-Meta is sketched in Barnden (1998) and a concurrent conference submission.

### Connections to Some Topical Issues

The research contributes to the debate between the Simulation Theory and Theory Theory of mental ascription (see, e.g., Carruthers & Smith, 1996), in that it clarifies what is realistically needed for both of these approaches to work, especially when uncertainty handling is involved. Most of the debate has been rather vague about the necessary underlying processes, and has hardly devoted any detailed attention to

uncertainty handling. The ATT-Meta research also extends simulation to apply to metaphorical reasoning, as a type of pretence.

ATT-Meta's metaphorical pretence processing appears to provide a partial implementation of the "conceptual blending" notion of Turner & Fauconnier (1995). Metaphorical pretence cocoons can contain a mixture of pretence-based and non-pretence-based reasoning (cf. the fact importation in a previous section). Uncertainty-handling, including conflict resolution, is something that conceptual blending needs in order to become algorithmically specific and conceptually plausible.

Some psychological research suggests the people may not construct literal meanings for metaphorical utterances, at least if the utterances are in an appropriate context and are of a familiar nature (see Gineste & Scart-Lhomme, in press, for a recent review). Although ATT-Meta is not meant to be a psychological model, it is worth noting that its use of literal meanings does not conflict with the psychological research. For one thing, the ATT-Meta research is not against the use of direct mappings from familiar *manifestations* of metaphors to metaphorical meanings (rather than from novel manifestations of familiar *metaphors*). Also, the psychological studies alluded to do not involve very detailed views of sentence processing, literal or otherwise, so that there is plenty of room to dispute that they indicate that literal meanings are not being constructed. For instance, an observation that it does not take longer to process a given utterance when taken metaphorically than when taken literally might indicate only the following: that the discourse-understanding *beyond* mere sentence-meaning construction (such as work for discourse-coherence establishment) in the literal case is replaced by analogous but different work in the metaphorical case, with the literal meaning itself being constructed in both cases and taking but a minor fraction of the overall reaction time. A further argument is given in Barnden (in press).

### Conclusion

The ATT-Meta-based research seeks to integrate the treatment of mental state reasoning, metaphorical reasoning and uncertainty. It thereby contributes to various areas with Cognitive Science and brings together concerns that have been artificially separated in the past. By not insisting on metaphorical utterances having meanings other than their literal ones, it has the freedom to deal powerfully with novel manifestations of (familiar) metaphors. Its pretence-based approach makes its handling of metaphor-based reasoning akin to its handling of belief reasoning.

### Acknowledgment

The research was supported in part by grant number IRI-9101354 from the National Science Foundation (USA).

### References

- Asher, N. & Lascarides, A. (1994). Intentions and information in discourse. In *Procs. 32nd Annual Meeting of the Association for Computational Linguistics*. New York: ACM.
- Ballim, A. & Wilks, Y. (1991). *Artificial believers: The ascription of belief*. Hillsdale, N.J.: Lawrence Erlbaum.
- Barnden, J.A. (1998). Uncertain reasoning about agents' beliefs and reasoning. Technical Report CSRP-98-11, School of Computer Science, The University of Birmingham, U.K. Invited submission to *Artificial Intelligence and Law*.
- Barnden, J.A. (in press). An AI system for metaphorical reasoning about mental states in discourse. In Koenig, J-P. (Ed.), *Conceptual Structure, Discourse, and Language II*. Stanford, CA: CSLI/Cambridge Univ. Press.
- Barnden, J.A., Helmreich, S., Iverson, E. & Stein, G.C. (1994). Combining simulative and metaphor-based reasoning about beliefs. In *Procs. 16th Annual Conference of the Cognitive Science Society*. Lawrence Erlbaum.
- 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), pp.101-123.
- Carruthers, P. & Smith, P.K. (Eds). (1996). *Theories of theories of mind*. Cambridge University Press.
- Chalupsky, H. (1996). Belief ascription by way of simulative reasoning. Ph.D. Diss., Dept of Computer Science, State University of New York at Buffalo.
- Cravo, M.R. & Martins, J.P. (1993). SNePSwD: A newcomer to the SNePS family. *J. Experimental and Theoretical Artificial Intelligence*, 5 (2&3), pp.135-148.
- Davidson, D. (1979). What metaphors mean. In S. Sacks (Ed.), *On Metaphor*. University of Chicago Press.
- Davies, M & Stone, T. (Eds) (1995). *Mental simulation: evaluations and applications*. Oxford, U.K.: Blackwell.
- Dinsmore, J. (1991). *Partitioned representations: a study in mental representation, language processing and linguistic structure*. Dordrecht: Kluwer Academic Publishers.
- Dragoni, A.F. & Puliti, P. (1994). Mental states recognition from speech acts through abduction. In *Procs. 11th European Conf. on AI*. Chichester, UK: Wiley.
- Gineste, M-D. & Scart-Lhomme, V. (in press). Comment comprenons-nous les métaphores? *L'Année Psychologique*, 1998.
- Haas, A.R. (1986). A syntactic theory of belief and action. *Artificial Intelligence*, 28, 245-292.
- Hobbs, J.R. (1990). *Literature and cognition*. CSLI Lecture Notes, No. 21, Center for the Study of Language and Information, Stanford University.
- Lakoff, G. (1993). The contemporary theory of metaphor. In A. Ortony (Ed.), *Metaphor and Thought*, 2nd edition. Cambridge University Press.
- Lakoff, G. & Turner, M. (1989). *More than cool reason: a field guide to poetic metaphor*. U. Chicago Press.
- Parsons, S., Sierra, C. & Jennings, N. (1998). Multi-context argumentative agents. In *Working Papers of the Fourth Symp. on Logical Formalizations of Commonsense Reasoning*, London, 7-9 January 1998.
- Turner, M. & Fauconnier, G. (1995). Conceptual integration and formal expression. *Metaphor and Symbolic Activity*, 10 (3), pp.183-204.