

An Implemented Context System that Combines Belief Reasoning, Metaphor-Based Reasoning and Uncertainty Handling

John A. Barnden and Mark G. Lee

School of Computer Science, The University of Birmingham
Birmingham, B15 2TT, United Kingdom
{J.A.Barnden,M.G.Lee}@cs.bham.ac.uk

Abstract. An implemented context-based reasoning system called ATT-Meta is sketched. The system can perform both reasoning about beliefs of agents and metaphor-based reasoning. In particular, it can perform metaphor-based reasoning about beliefs and reasoning acts. The metaphor-based reasoning and belief reasoning facilities are fully integrated into a general framework for uncertain reasoning. This framework allows for uncertain reasoning and conflict resolution within individual contexts, uncertainty in individual inter-context bridging rules, conflict resolution between the effects of different bridging rules, and conflict-resolution across context boundaries when argumentation inside a context conflicts with argumentation outside.

1 Introduction

We have developed and implemented a context-based mechanism for belief reasoning, in our ATT-Meta system (Barnden 1998, Barnden *et al.* 1994), mainly for the purposes of natural language pragmatics. However, we have also taken the unusual step of developing, within the very same system, a context-based mechanism for metaphor-based reasoning—the reasoning needed during the understanding of metaphorical utterances. (“ATT-Meta” = “[propositional] ATTitudes” and “Metaphor.”) Use of contexts for handling metaphor is a central part of the work of Fauconnier & Turner (1998). However, that work has not led to an implemented metaphor-processing system, does not combine belief reasoning and metaphor-based reasoning, and does not address uncertainty issues in any detail.

Indeed, belief reasoning and metaphor are almost always studied separately. But metaphorical utterances can be about the mental states and processes of agents. ATT-Meta was originally focused only on that type of metaphor, though in fact its metaphor mechanisms are in no way confined to it. Metaphorical utterances about mental states lead to the need to combine belief reasoning and metaphor-based reasoning in intimate ways, one of which we will detail.

A third prominent strand in our work is qualitative uncertainty-handling. This is important both for useful, common-sensical belief reasoning and useful, common-sense reasoning based on metaphor. Although belief reasoning is occasionally combined with uncertainty handling (e.g., in Parsons, Sierra & Jennings 1998), the only previous metaphor research that devotes extensive technical attention to uncertainty appears to be that of

Hobbs (1990). One goal of the ATT-Meta research is for the system to perform uncertain belief reasoning and uncertain metaphor-based reasoning in a systematic and uniform way, doing more justice to uncertainty than has heretofore been seen.

Although we have only studied contexts for belief reasoning and metaphor-based reasoning, our work on uncertainty should be generalizable to contexts much more generally. The mechanisms we have developed allow for uncertain reasoning and conflict resolution within individual contexts, uncertainty in individual inter-context bridging rules, conflict resolution between the effects of different bridging rules, and conflict-resolution across context boundaries (e.g. to handle conflict between what it seems an agent believes according to simulation of the agent's reasoning, and what it seems the agent believes according to "external" information, such as statements by other people about the agent). The mechanisms have been designed to work with unlimited nesting of contexts, where contexts of different types can occur in a nesting.

The remainder of the paper addresses the following topics, in sequence. the type of metaphorical utterance we address; ATT-Meta's basic reasoning facilities and uncertainty-handling; ATT-Meta's facilities for reasoning about agents' beliefs and reasoning; ATT-Meta's metaphorical reasoning; and one type of interaction of belief reasoning and metaphor-based reasoning. Our presentation will be informal, as it is the general issues that our work has addressed and the inference strategies that we have used that are important, not their precise clothing in formal terms.

2 Metaphor in ATT-Meta

A metaphorical utterance is one that *manifests* a metaphorical view, where a *metaphorical view* is a conceptual view of one topic or domain as another that is judged qualitatively different from it. Here we broadly follow Lakoff (e.g., Lakoff 1993). An example of a metaphorical view is the view of MIND AS PHYSICAL SPACE. A natural-language manifestation of this view might be "John believed in the recesses of his mind that" In a manifestation, the topic actually being discussed (John's mind, in the example) is the *target* topic or domain, and the topic that it is metaphorically cast as (physical space, in the example) is the *source* topic or domain.

The ATT-Meta system does not currently deal with novel metaphorical *views* that can appear in natural language utterances. Rather, it has pre-established knowledge of a specific set of metaphorical views, including MIND AS PHYSICAL SPACE. But it is specifically designed to handle novel *manifestations* of those views, or, more precisely and generally, manifestations which cannot immediately be handled by the known mappings from source to target that constitute the metaphorical view. Such manifestations are common in ordinary discourse (see, e.g., Martin 1994).

Mundane types of discourse, such as ordinary conversations and newspaper articles, often use metaphor in talking about mental states/processes of agents. There are many mental-state metaphorical views 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;" 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 "One part of him was convinced that he should go to the party."

3 ATT-Meta's Basic Reasoning

ATT-Meta is a rule-based reasoning system. Reasoning (sub)goals, provided facts and derived propositions are all called *hypotheses* for convenience. Hypotheses are actually (Barnden *et al.* 1994) terms in a first-order, episode-based logic akin to that of, say, Hobbs (1990), but for simplicity of presentation they will be portrayed here as formulas in a first-order modal logic.

For our present purposes a fact contains no variables, and if a goal contains variables then the only quantification is implicit existential quantification of all those variables, prefacing the entire formula. A fact or goal can involve negation, conjunction and disjunction, but in this paper we will only touch on negation. The reader may assume that for each ground (i.e., variable-free) hypothesis entertained, the *complement* is also entertained, even though this is only roughly true. Two variable-free hypotheses are complementary when one is the negation of the other. This paper largely omits explicit consideration of variables in hypotheses from now on.

At any time any particular hypothesis is tagged with a *qualitative certainty level*, one of certain, presumed, suggested, possible or certainly-not. The intuitive meanings are as follows:

certain: The hypothesis is true without reservation.
certainly-not: The complement is certain.
presumed: the hypothesis is a *default*: i.e., it is taken as a working assumption, pending further evidence.
suggested: there is evidence for the hypothesis, but it is not (yet) strong enough to enable the hypothesis to be presumed.
possible: No evidence at all has yet been found for the hypothesis, but the negation is not certain.

(The level *certainly-not* is included for practical reasons, even though it is redundant.) If a hypothesis is at a level less than *certain* at some moment it can potentially achieve a higher level later. However, each hypothesis also has an upper bound that can limit the rise of its certainty level. Our current system of levels is no doubt simplistic compared to what may ultimately be required, but it is complex and powerful enough for substantial advances to be made in our areas of interest.

A rule in ATT-Meta is of largely traditional form, having an *antecedent* and a *consequent*, but it also has a *qualitative certainty level*. We will display rules in the following format:

```
IF <antecedent component 1> AND <antecedent component 2> . . . . AND <antecedent component N>  
THEN [ <certainty level> ] <consequent>
```

The consequent is of the same form as a goal, and may therefore contain variables. Each component of the antecedent is an expression in goal form. Implicitly, all variables in the rule are universally quantified, outside the whole rule. The certainty level of a rule is one of *suggested*, *presumed*, *certain*.

ATT-Meta applies its rules in a backchaining style. It is given a reasoning goal by a user, and uses rules to generate subgoals in the normal way. Goals can of course also be satisfied by provided facts. When a rule application successfully supports a goal, it supplies a level of certainty to it, calculated as the minimum of the rule's own certainty level and the levels picked up from the hypotheses satisfying the rule's antecedent.

The rule's certainty level thus acts as an upper bound on the certainty that the rule can supply to its conclusion. The rule only supplies support to the goal in this way if all those certainty levels are at least `suggested` (because every hypothesis starts off as `possible` anyway). When several rules support a hypothesis, the maximum of their certainty contributions is taken. Currently, ATT-Meta tries to apply all rules to all goals, though an important future research goal is to refine this procedure by the use of heuristics.

The system contains a modified type of justification-based truth-maintenance mechanism for propagating levels of certainty around via rule-applications. There are often cycles in the *hypothesis graph* (which is essentially a dynamically changing AND/OR graph defined by the rule-applications that have been created). As a result, in general the system iteratively settles to a consistent set of hypothesis certainty levels. The process is guaranteed to terminate.

4 Basic Conflict Resolution

When both a hypothesis and its complement are supported to level at least `presumed`, conflict-resolution will take place. The most interesting case is when both hypotheses are exactly `presumed`. In this case, the system attempts to see whether one hypothesis has *more specific* evidence than the other. The former is then the "winner." If the system determines a winner, it downgrades the certainty level of the loser to `suggested`, keeping the winner's at `presumed`. If it cannot determine a winner, it downgrades both to `suggested`. It is therefore conservative in its conflict resolution, avoiding arbitrary decisions between hypotheses. For instance, in the Nixon diamond it will end up attaching level `suggested` to the hypotheses that Nixon is and is not a pacifist.

Specificity comparison is a commonly used heuristic for conflict-resolution in AI (see, e.g., Loui *et al.* 1993) although serious problems remain in coming up with adequate and practical heuristics. ATT-Meta's specificity comparison depends on what facts the two conflicting hypotheses rely on and on derivability relationships between the hypotheses supporting the conflicting hypotheses. The derivability aspect is analogous to derivability aspects of other schemes for specificity, and has as a special case the ancestor-closeness heuristic central to most semantic networks. The fact aspect relies on the intuition that if one of the conflicting hypotheses relies on more facts than the other then it has more specific evidence ("more" in a set-inclusion sense).

The following is a simple example of specificity comparison is sufficient for the purposes of this paper. It illustrates only the derivability aspect. Suppose the system has rules saying that students are presumably not rich, medical students are presumably rich, but English medical students are presumably not rich. Thus the second rule overturns the first and the third overturns the second. Suppose there is fact saying that Bill is an English medical student. Then the system concludes that Bill is presumably not rich, defeating the line of evidence from the second rule mentioned. The system can deal with any number of levels of overturning in a scenario of this sort.

5 ATT-Meta's Reasoning about Agents' Beliefs and Reasoning

ATT-Meta has facilities for reasoning non-metaphorically about the beliefs and reasoning acts of agents, where those beliefs and acts may in turn be about the beliefs and reasoning for further agents, and so forth. Although ATT-Meta can reason about beliefs

in an ordinary rule-based way (e.g. using rules such as “if X is English then X presumably believes cricket is interesting”), its main tool is *simulative reasoning* (e.g., Creary 1979). Roughly, in attempting to show that agent X believes P, the system puts P as a goal in a *simulation context* for X, which is a special environment which is meant to reflect X’s own reasoning processes. If P is supported within the simulation context by hypotheses Q, where the system takes X to believe the Q, then the reasoning from the Q to P in the context is alleged (by default) to be reasoning by X. The reasoning within the context can involve ordinary rule-based reasoning and/or simulation of further agents. Currently, ATT-Meta can use any of its rules within any simulation context, so that all rules are effectively treated as common knowledge. Of course, this needs to be modified in future research.

The picture just given of simulative reasoning needs to be greatly complicated because of uncertainties. First, the reasoning within the simulation context can be uncertain. Also, the information about X resulting from a simulation of X is itself uncertain: even if the reasoning *within* the simulation context supports P with absolute certainty, ordinary rule-based reasoning outside the context may support the proposition that it is *not* the case that X believes P. Also, there is always the possibility that (a) X may not in fact have done some of the alleged reasoning steps, and (b) even if X has done them, it may be that X has also done steps that support the negation of P (when P is uncertain within the context).

An *atomic belief hypothesis* is a formula of the form $B(a, \phi, \lambda)$, where B is a modal operator, a is a term denoting the agent, ϕ is a hypothesis-like formula, and λ is a certainty level greater than or equal to possible. Intuitively the hypothesis means that the agent believes ϕ with certainty AT LEAST λ . Because of the λ argument, the system expresses uncertainty in every layer of belief. Hypotheses involving belief take part in rule-based reasoning just as other hypotheses do (subject to a rather complicated optimization concerning the λ arguments in belief hypotheses). Notice that because of the “AT LEAST λ ” provision, if $B(a, \phi, \lambda)$ holds then so does every $B(a, \phi, \lambda')$ for lesser certainty levels λ .

Ordinary rule-based reasoning involving belief hypotheses is not intended to cater for reasoning that is alleged to be performed by other agents. Simulative reasoning is implemented for this purpose. Whenever a hypothesis that is an atomic belief hypothesis is created, the hypothesis formed by stripping off the application of the belief operator B is also formed (unless it already exists), and placed inside the “simulation context,” or “reasoning space,” for the agent in question (similar to the belief spaces and contexts in many other authors’ work). This happens recursively in the case of nested belief hypotheses, so that we get nested reasoning spaces (nested contexts). Conversely, whenever a goal is created inside an agent’s reasoning context, an application of B is wrapped around it and placed in the immediately surrounding context (which may be the system’s own).

The support for an atomic belief hypothesis can come from: ordinary rule-based reasoning in the layer (i.e., reasoning context) it lies in; from simulative reasoning resulting from stripping off the B; and from rule-based reasoning at higher levels resulting from adding further B applications. One of our foci in this paper is on matters of conflict resolution, discussed below. For this purpose, a technicality needs to be explained. Each rule-application within a non-system reasoning context is “lifted” to the next reasoning-context out. That is, consider a hypothesis $B(a, \phi, \lambda)$ within a reasoning context R , and hypothesis ϕ within R_a , the reasoning context for a nested within R . For each rule-application α supporting ϕ , a “lifted rule application” $A(\alpha)$ is created as a support for $B(a, \phi, \lambda)$. Let $\Delta(\alpha)$ be the set of hypotheses directly used by application α .

Then, the set of hypotheses directly used by $\Lambda(\alpha)$ formed by wrapping $B(a, \dots)$ around each member of $\Delta(\alpha)$, plus an extra “agent-inference hypothesis” that will be explained in the discussion of metaphor. The system always uses `presumed` as the strength of the imaginary rule of which $\Lambda(\alpha)$ is an application.

Such lifted rule applications account for almost all the direct effect that the simulation of a (i.e., the reasoning within R_a) has on the reasoning about $B(a, \phi, \lambda)$ within context R . Importantly, *lifted rule applications for a hypothesis are treated in the same way as ordinary rule applications for the hypothesis in respect of certainty-value accumulation and conflict resolution involving the hypothesis*. This principle supplies a valuable amount of modularity in the design of the system.

6 Conflict Within and Across Belief Layers

Consider the hypothesis $B(a, \phi, \text{presumed})$ within a reasoning context R , and hypothesis ϕ within R_a . Also consider the following related hypotheses: the complement ϕ' of ϕ within R_a , the hypothesis $B(a, \phi', \text{presumed})$, and the complements of the two B hypotheses. (For simplicity, we will not discuss here the matter of other B hypotheses of the form $B(a, \phi, \dots)$). All six hypotheses can have independent rule-based argumentation supporting them, and we are therefore faced with the problem of there being potential conflict-resolution within R_a between ϕ and ϕ' , and within R between $B(a, \phi, \text{presumed})$ and its complement, and between $B(a, \phi', \text{presumed})$ and its complement. In two special cases it is clear what to do.

(a): when there is no evidence for the four hypotheses in R other than what arises within R_a , conflict resolution should occur within R_a , between the inside hypotheses ϕ and ϕ' . As a result, the certainty levels of some or all of the four outside hypotheses, as well as of one or both of the inside hypotheses, will be adjusted.

(b): This special case is the opposite of case (a) — that is, no rule-based reasoning inside R_a supports ϕ or ϕ' to level `presumed`. Then, conflict resolution should not be done within the simulation, but should occur outside.

The difficult case is when there is (non-lifted) rule-based support, either directly or indirectly, for at least one of the four outer hypotheses, as well as rule-based support (lifted or otherwise) for at least one of ϕ and ϕ' . This case is made to look like special case (b) by means of the lifting mechanism. That is, the rule-based argumentation within context R_a is lifted to the outer context, and combines with the rule-based argumentation in that context. Thus, conflict-resolution is performed in the normal way in that context, in an attempt to find a winner between $B(a, \phi, \text{presumed})$ and its complement, and between $B(a, \phi', \text{presumed})$ and its complement.

However, a special provision helps this process. Any rule-applications directly supporting $B(a, \phi', \text{presumed})$ are copied to become rule applications directly supporting the complement of $B(a, \phi, \text{presumed})$ as well. Conversely, applications directly supporting $B(a, \phi, \text{presumed})$ are copied to become ones for the complement of $B(a, \phi', \text{presumed})$ as well. This is necessary so as to allow conflict within the simulation to be reflected as conflict in outer contexts, because $B(a, \phi', \text{presumed})$ and $B(a, \phi, \text{presumed})$ are not complements of each other, and because specificity-comparison depends in part on *direct* support relationships between hypotheses.

Since the cross-context considerations we have been discussing can apply at any context boundary in a nesting of reasoning contexts, the following recursive process

is defined. Starting at the top reasoning context (i.e., the system's), the system works down through contexts. If it hits an atomic belief hypothesis and its complement that are in conflict, and where at least one has non-lifted rule-based support (either directly or indirectly), it applies the basic conflict-resolution procedure to those hypotheses. The system now descends to the inner context, and so forth. If the conflicting belief hypotheses do not have non-lifted rule-based support, then the system just goes down to the inner context.

7 Metaphor-Based Reasoning in ATT-Meta

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. The system does reasoning on the basis of these formulae. Consider the following metaphorical sentence:

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

The sentence manifests two metaphorical views: MIND AS PHYSICAL SPACE and IDEAS AS PHYSICAL OBJECTS. We assume that the understander is familiar with both metaphorical views. We assume that the understander only has a physical sense for “store-room” and that the notion of physical store-room is not *map-enshrined* for the understander by either of the metaphorical views: that is, the notion is not mapped by either metaphorical view to any target-domain notion. We say the sentence is an *exploitation* of the two metaphorical views, because of this use of a non-map-enshrined notion. Such exploitative sentences are a major issue, because they appear to be common in discourse. However, existing implemented system for metaphor other than our own are largely not directed at them, with two main exceptions: Martin (1994) and Narayanan (1997). These systems are more restricted in scope than ours specialised aims than ours and do not deal with the uncertainty-handling issues that we address.

In our approach to exploitative metaphorical utterances, we try to avoid discovering or inventing mappings for non-map-enshrined 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* map-enshrined 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.

The desirability of deriving the connotation could arise from a belief-reasoning episode, but we defer this matter to Section 3.8. In the present section we 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 literally*, i.e. taken to say that John’s mind literally 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.

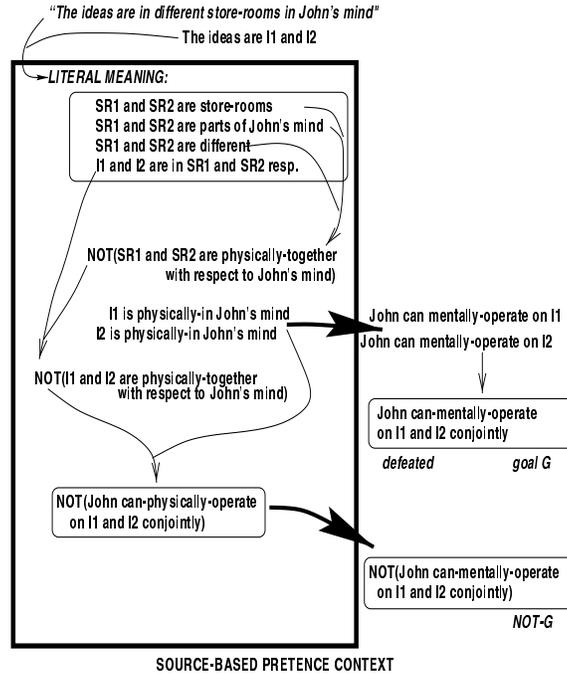


Fig. 1. Part of the reasoning in store-room example. The literal meaning of the sentence is shown in the topmost part of the large box. The processing shown within the large box is source-based inference. The processing outside the box is reasoning within the terms of the target domain. The arrows from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

The reasoning, apart from the application of the conversion rules, occurs within a special context that we call a *source-based pretence context*. (Elsewhere we have used the word “cocoon” instead of “context.”) Within this context, the *literal* information from the sentence—(A) above—is taken to be true, no matter how absurd. This meaning, being literal, is in the terms of the source domains of the metaphorical views used in the sentence, that is the domains of PHYSICAL SPACE and PHYSICAL OBJECTS in our example. ATT-Meta can use any of its reasoning rules within the source-based pretence context 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 context. The reasoning within the pre-

tence context 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 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 "operation" 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 [presumed] 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 context contains H and that the level of certain for H within that context is at least presumed. A hypothesis of form "it is being pretended that, with certainty at least λ , H" is a highly analogous to a belief hypothesis. Using the same style of notation as we used above for belief hypotheses, it can be expressed as a formula of the form $\text{pretend}(H, \lambda)$. Such a formula is outside the pretence context. The formula H is within the context. A formula using the pretend operator is called a "pretence hypothesis." To oversimplify this, every such hypothesis is reflected within the context by its "H" argument, and every hypothesis H within the context is reflected outside by a pretence hypothesis.

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 converse-contrapositive 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 {I1,I2}), where I1 and I2 are the ideas mentioned in the sentence, from the premise that I1 and I2 are literally in different store-rooms in John's mind. This is because (a) store-rooms 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 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 context. Into that context is inserted (at the appropriate moment in the backwards rule-based reasoning) the facts that I1 is in a store-room SR1, I2 is in a store-room SR2, SR1 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 context, that John’s mind is a building and therefore a physical region, and I1 and I2 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 *literal* meaning of the utterance, conducted within the pretence context, 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 somewhat artificial assume the goal is

(G) John can-mentally-operate-on {I1, I2}.

In fact, the system will conclude with a certainty of *presumed* for the *negation* of (G), thereby drawing the Connotation displayed above. Below we will show how the need to establish the Connotation could arise naturally and implicitly in belief reasoning.

8 A Case of Metaphorical Conflict

Notice the uncertainty in the system’s conclusion about the negation of (G)—it is only *presumed*. The uncertainty arises partly from the fact that conversion rules, such as (2), have a certainty level only of *presumed*, not *certain*. Indeed, ATT-Meta allows the possibility that there is evidence for (G) itself that outweighs the evidence for its negation. In our actual implementation of the store-rooms example, we do have a further rule, operating entirely within the target domain of mental states:

(3) IF X can-mentally-operate-on J AND X can-mentally-operate-on K
THEN [*presumed*] X can-mentally-operate-on {J,K}.

This supports (G), because of a MIND AS PHYSICAL SPACE conversion rule that says that if an idea J is in a mind that is viewed as physical region then the agent can mentally operate on the idea. Thus, there is a real conflict between (G) and its negation.

Our general-purpose, specificity-based conflict-resolution mechanism is adequate for establishing that the evidence supporting goal (NOT-G) above it be more specific than the evidence supporting (G) itself. The argumentation for the negation of (G) ultimately relies on all the facts that (G) relies on together with an extra one, namely that it

is pretended that SR1 and SR2 are *different*. Therefore, merely by the fact-based aspect of specificity comparison, (NOT-G) is regarded as having more specific support.

(G) and (NOT-G) are both supported directly or indirectly by conversion rules (for different metaphorical views). However, (G) could potentially be supported by information that is entirely derived in the target domain. For instance, the understander might know a fact F from previous non-metaphorical discourse that allows it to be concluded that John is entertaining I1 and I2, and therefore can mentally operate on I1 and I2 individually, so that by target-domain rule (3) we get further support for (G). However, although we do not show here how this happens, (NOT-G) would still win against (G), again just by the *general-purpose* conflict-resolution mechanism.

9 Uncertainty in Metaphor

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

(1) Given an utterance, it is often not certain what particular metaphorical views or variants of them are manifested. Correspondingly, ATT-Meta may merely have *presumed*, for instance, as a (tentative) level of certainty for a pretence hypothesis, such as that the pretence that idea I1 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.

(2) Knowledge about the source domain of the metaphorical view is itself generally uncertain. Correspondingly, in ATT-Meta the hypotheses and reasoning within the pretence context are usually uncertain, and can involve the use of the conflict-resolution mechanism described above. For instance, it 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.

(3) 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 specificity 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. However, many researchers (e.g., Lakoff 1993) appear to assume that, in cases of conflict, target information should override metaphor-based inferences, and thus do not fully address the potential uncertainty of target 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 target domain. To metaphorically 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.

(4) Different conversion rules working may create propositions that lead to conflict. We saw this in the discussion of (G) and its negation in the store-room example. Such conflict can also occur between the conversion rules for just a single metaphorical view. For example, in uses in real discourse of the metaphorical view of ORGANIZATIONS AS PHYSICAL OBJECTS, companies are often cast as eating things, such as other companies. Eating actions by companies might be mapped to *legal* company-acquisition

be default, but there could be a more specific mapping that says that eating by a Mafia-run company maps (by default) to *illegal* company acquisition.

10 Importation-Based Conflicts

By default, a pretence context is taken to contain as a fact any fact sitting immediately outside. The store-room example provides an illustration of the need for this *importation* of facts. One of the rules applied within the pretence context needs to establish that John is a person. That John is a person is given only outside the context, however.

Importation can lead to reasoning conflicts. In brief, within the pretence context, the imported facts may support the complement of a `presumed` hypothesis that derives indirectly from the special metaphorical facts inserted into the context at the start (e.g., the fact that idea I1 is in store-room SR1). That fact leads to the conclusion that I1 is a physical object. But the fact outside the context that I1 is an idea leads by a rule about mental entities that I1 is *not* a physical object. To cater for such situations, *certain* facts outside the pretence context are downgraded to `presumed` on importation, and a simple modification is made to the conflict-resolution scheme as described so far: within a metaphorical pretence context, 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-comparison is attempted again. This preference against imports needs to be reflected out to higher levels via lifting. The preference can naturally be regarded as an additional specificity principle.

11 Combining Metaphorical and Belief Reasoning

Here we present one type of interaction that is possible between metaphorical and belief reasoning in ATT-Meta. Other types not addressed here include reasoning about other agents' metaphorical reasoning and reasoning about the alleged reasoning of personified non-agents.

The store-room example as presented so far is does not address the question of why understanders would try to establish the above Connotation, or how understanders would make use of that connotation if they spontaneously derived it. But suppose that the two ideas (I1, I2) mentioned in the store-rooms sentence (1) are stated (elsewhere in the discourse) to be believed by the agent, John, and suppose that they lead by some fairly direct, natural reasoning to some conclusion C. Then, if (1) had not been uttered, it would have been reasonable to ascribe by default to John a belief in the conclusion C. But this ascription relies on a default assumption that John does mentally operate conjointly on I1 and I2. If there is evidence that this is not the case, then the ascription should be abandoned. Sentence (1) provides just such evidence against John's mentally operating conjointly on I1 and I2. We just mention briefly here what happens, with the aid of Fig. 2. We suppose the user supplies the goal

(K) John believes C to level at least `presumed`

and also the facts that John believes I1 to level `certain` and that John believes I2 to level `certain`. Without any further information, ATT-Meta would proceed as follows, on the simplifying assumption that C can be derived to level `presumed` by just

one rule application from I1 and I2. ATT-Meta would open a simulation pretence context in which I1 and I2 are set up as facts and C is set up as a goal. An attempt would be made to establish C within the simulation context. We assume that this attempt is successful because of the mentioned rule application, giving some certainly level for C. For definiteness, let us say that this level is *certain*. As a result, the hypothesis (K) in the top context is supported. The level to which it is currently supported is also *presumed*, according to a general policy regarding simulation of reasoning.

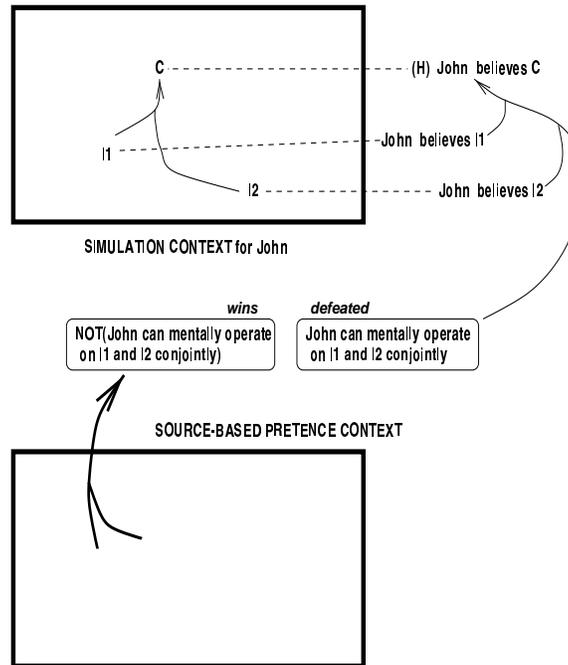


Fig. 2. Combining belief reasoning and metaphorical reasoning.

But, the system notes that hypothesis (K) depends not only on John's belief in I1 and I2 but also on the assumption that John mentally operates on I1 and I2 conjointly. Such notes are established in all cases of belief reasoning, not just when metaphorical reasoning is in the offing: there may, in principle, be non-metaphorical evidence that the assumption is wrong. Because of the note, the assumption is set up as a goal, namely goal (G) of the previous section. Such goals are the "agent-inference hypotheses" mentioned in an earlier section. They are given an initial certainty level of *presumed*, and this level stays if there is no counter-evidence. In our example, the evidence for (NOT-G) prevails. Because (G) is defeated, the support for the hypothesis that John believes C to at least level *presumed* collapses. (And as a result, the inference step within the simulation pretence context is deleted.)

12 Conclusion

ATT-Meta embodies a unique, implemented application of context-handling. It incorporates belief reasoning, metaphor-based reasoning and uncertainty-handling in a unified framework. In addition, the uncertainty-handling could be generalized to work with types of context other than those needed for belief reasoning and metaphor. We have paid much attention to the need for conflict resolution to be handled correctly within different contexts at any depth of nesting, and across the boundaries of contexts. Also, different bridging rules working from one context to another may lead to conflict. We saw a case of this in metaphor-based reasoning.

ATT-Meta's metaphorical reasoning in all its aspects (reasoning within contexts, application of conversion rules, etc.) uses the same general-purpose rule-based reasoning provisions as for any other reasoning within the system. The only thing special about it is the type of knowledge used (mainly in the form of conversion rules), fine detail of the handling of the contexts, and the special conflict-resolution provisions concerning importation of facts into metaphorical pretence contexts.

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 "issues" raised by surrounding sentences.

References

1. Barnden, J.A. (1998). Uncertain reasoning about agents' beliefs and reasoning. Technical Report CSR-98-11, School of Computer Science, The University of Birmingham, U.K. Also in *Artificial Intelligence and Law*, in press.
2. Barnden, J.A., Helmreich, S., Iverson, E. & Stein, G.C. (1994). An integrated implementation of simulative, uncertain and metaphorical reasoning about mental states. In J. Doyle, E. Sandewall & P. Torasso (Eds), *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourth International Conference*, pp.27–38. San Mateo, CA: Morgan Kaufmann.
3. Creary, L. G. (1979). Propositional attitudes: Fregean representation and simulative reasoning. *Procs. 6th. Int. Joint Conf. on Artificial Intelligence* (Tokyo), pp.176–181. Kaufmann.
4. Fauconnier, G. & Turner, M. (1998). Conceptual integration networks. *Cognitive Science*, 22(2), pp.133–187.
5. Hobbs, J.R. (1990). *Literature and cognition*. CSLI Lecture Notes, No. 21, Center for the Study of Language and Information, Stanford University.
6. 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.
7. Loui, R.P., Norman, J., Olson, J. & Merrill, A. (1993). A design for reasoning with policies, precedents, and rationales. In *Fourth International Conference on Artificial Intelligence and Law: Proceedings of the Conference*, pp.202–211. New York: ACM.
8. Martin, J. (1994). Metabank: A knowledge-base of metaphoric language conventions. *Computational Intelligence*, 10 (2), pp.134–149.
9. Narayanan, S. (1997). KARMA: Knowledge-based action representations for metaphor and aspect. Ph.D. thesis, EECS Department, U. of California, Berkeley, August 1997.
10. Parsons, S., Sierra, C. & Jennings, N. (1998). Multi-context argumentative agents. In *Working Papers of the 4th Symp. on Logical Formalizations of Commonsense Reasoning*, London.