

Uncertainty and Conflict Handling in the ATT-Meta Context-Based System for Metaphorical Reasoning

John A. Barnden

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

Abstract

At CONTEXT'99, the author described the ATT-Meta context-based system for (a) reasoning uncertainly about agents' beliefs and (b) performing some of the uncertain reasoning needed for the understanding of metaphorical language. ATT-Meta's handling of uncertainty is qualitative, and includes provisions for adjudicating conflicts between different lines of reasoning. But most of the detail on conflict-handling given in the earlier paper concerned conflicts arising for the special requirements of (a). Furthermore, there have been major recent changes in the conflict-handling approach. The present paper provides a detailed account of major parts of the current approach, stressing how it operates in metaphorical reasoning. In concentrating on uncertainty-handling, this paper does not seek to repeat the justifications given elsewhere for ATT-Meta's approach to metaphor.

Acknowledgment

The research is being supported by grant GR/M64208 from the Engineering and Physical Sciences Research Council of the U.K.

1 Introduction

Various types of uncertainty and conflict are important in the understanding of metaphorical utterances. They have crucially affected much of our detailed theoretical and implementational development of the ATT-Meta approach and system (see Barnden references, Barnden *et al.* references and Lee & Barnden 2001). As explained in Barnden & Lee (1999), some contexts in ATT-Meta serve as belief spaces within which to simulate alleged reasoning of agents, while others serve as spaces within which to conduct reasoning conforming to “pretences” set up by metaphorical utterances. For instance, in the case of

“In the far reaches of her mind, Anne believes that Kyle is unfaithful.”

(slightly edited version of a real-discourse example) the context is one in which Anne’s mind is taken actually to be a physical space that has far reaches. Some consequences of this are teased out by reasoning that is typically uncertain, and then operated on by mapping rules (context-bridging rules) to create hypotheses about Anne’s mind, outside the pretence context. This process should become clearer below. Thus, our approach is heavily context-based, especially as contexts can be nested within each other in several different ways (Barnden 1998). The uncertainty and conflict-handling included in our approach and system operates both within individual contexts and “across” contexts in the sense that flow of information across context boundaries can cause conflict with other information within contexts. It is plausible that our provisions can be adapted for use in context-based reasoning systems that have nothing to do with metaphor.

In our approach metaphor in general is the phenomenon of thinking or communicating about something X in a way that relies on or is motivated by seeing it as something qualitatively different from it. Thus, in the Anne/Kyle example X is Anne’s mind, which is being seen as a physical space of a sort that allows for “far reaches.” This metaphorical view (conceptual metaphor—Lakoff 1993) of MIND AS PHYSICAL SPACE is very commonly used in English natural language discourse, and we will assume that a normal user of English is accustomed to its use, and possesses a suitable mapping between aspects of physical space and aspects of mind. Nevertheless, we do not assume that in the Anne/Kyle example the understander is able to map the concept of “far reaches” of a physical region to something to do with the mind. Rather, the understander must be able to connect this concept to other source-domain concepts that *are* mapped. The connecting is the within-pretence-context reasoning mentioned above. We will also assume that the metaphoricity is not “sidelined” by a mind-domain lexical sense for the phrase “in the far reaches” that permits the understander to avoid mapping from the source domain to the target domain.

Some of the types of uncertainty and conflict that ATT-Meta handles, at least to a useful extent, are:

1. Inferential links within the target domain of a metaphorical view can be uncertain, and there can be conflict between different inference chains in that domain.
2. Similarly for the source domain.
3. Reasoning within source-domain terms sometimes needs to interact with facts imported from the target domain, and conflict can arise as a result.

Explanation of (3) will be deferred until Section 7. Concerning point (2), understanding of the type of metaphor addressed in this paper requires possibly-substantial reasoning in the terms of the source domain, as we will see below. Assuming that the source domain is a commonsense one (e.g., the commonsense domain of physical space and physical objects), reasoning within it will be full of uncertainty and conflict. Consider “One part of Mike was insisting that Sally was right, but another part believed she was wrong.” This rests on a metaphorical view of MIND PARTS AS PERSONS, in which the mind (or a person) is viewed as having parts that are sub-persons, able themselves to have mental states, communicate in natural language, etc. When someone insists something, they generally believe it. So the understander can infer, in source-domain terms, that the first “part” mentioned in the sentence probably believed that Sally was right. The reasoning process rests purely on general knowledge about verbal communication and related mental states, which are aspects of the source domain of the metaphor. As for possible conflict in the source domain, this arises simply because commonsense domains generally contain exceptions to general rules or principles.

As for (1), equally, when the target domain is a commonsense one, reasoning within it is likely to be uncertain and conflictful. This applies just as much to reasoning episodes in the target domain that are triggered by transfers of information from the source domain as to reasoning episodes arising independently in the target domain. Handling of the three types of uncertainty listed, and others, is seamlessly integrated in ATT-Meta.

2 Main Operations in Metaphor Understanding in ATT-Meta

Our approach encapsulates the proposed source-domain reasoning within a special computational context called a *pretence cocoon*. A pretence cocoon is akin to the more familiar idea of a belief space. In general there can be more than one pretence cocoon (e.g., because of metaphor compounding: Lee & Barnden 2001), but this article will assume just one. Use of the cocoon is tantamount to the understander pretending to take the source-domain scenario alluded to in the utterance as a true statement. The propositions inside the pretence cocoon can be thought of as forming a description of an imaginary, pretend world. Propositions outside are about reality as the understander sees it, and will include, but not be restricted to, propositions about the target domain of the metaphorical utterance. Metaphorical mapping acts operate between the inside of the pretence cocoon and outside.

Thus, for the Anne/Kyle example, the reasoner pretends, within the cocoon, that Anne’s mind *really is* a physical space and that the believing really does occur in the far reaches of this space. Consequences of this, using knowledge about physical space, can be inferred in the pretence cocoon, possibly by substantial amounts of reasoning, and *some* consequences may (hopefully) then be able to be transformed, via mapping rules forming part of the relevant metaphorical views, to become propositions about the Anne’s state of mind that are outside the cocoon. We will call the reasoning within the cocoon “within-pretence reasoning.” Within-pretence reasoning is similar to the *simulative reasoning* used or proposed by many researchers for reasoning about agents’ beliefs and reasoning (see, e.g., Creary 1979, Davies & Stone 1995).

In the Anne/Kyle example, within-pretence reasoning could infer that the idea that Kyle is unfaithful cannot be physically operated upon (more than to a very low degree) by the Anne's conscious self (metaphorically viewed as a person within the mind-region), because that self is in a *central* part of the mind-region, not in the far reaches. This proposition is, we assume, part of the understander's knowledge of the MIND AS PHYSICAL SPACE metaphorical view. Then, if part of the understander's knowledge of the IDEAS AS PHYSICAL OBJECTS metaphorical view is that ability to engage in (pretended) physical operation on an idea corresponds to ability to engage in (real) conscious mental operation on the idea, then the understander can infer that Anne cannot consciously operate, mentally, on the idea that Kyle is unfaithful. More detail of the reasoning performed by ATT-Meta on the example is given in Figure 1 (at end of paper) and below.

We can summarize the main types of processing activity that we claim should take place in understanding of the type of metaphorical utterances of interest as follows:

- construction of the direct meaning of the utterance, i.e. the meaning it has by taking only the source-domain senses of the metaphorical words/phrases in the utterance.¹
- placing of the direct meaning in the *pretence cocoon*, which is a special computational context
- usually, performance of (predominantly) source-domain reasoning within the pretence cocoon, using that direct meaning together with general knowledge about the source domain
- possibly, application of *ancillary assumptions* associated with the relevant metaphorical views to create further propositions within the pretence cocoon (an example of this is the assumption in the Anne/Kyle discussion above that Anne's conscious self is in the central part of the mind-space)
- source-to-target or target-to-source mapping acts, including at least one that goes from source to target (e.g. mapping of inability to operate on an idea physically to inability to consciously, mentally operate on it).

This listing does not imply any particular temporal ordering of the types of processing. Indeed in ATT-Meta the reasoning works backwards from queries and can involve any intertwining of the above types of reasoning.

3 Irreality of Domains

Any non-fact rules pertaining to any domain (including even the target domain) can be considered for application inside the pretence cocoon. One reason for allowing this, despite the fact that the pretence cocoon is mainly for conducting reasoning in source-domain terms, is that the source domain and target domain in a metaphor can overlap. The possibility of overlap is not often discussed in the literature, but is pointed out by Kittay (1989) for example, and we see in the MIND PARTS AS PERSONS example in

¹ This is a highly simplified statement, ignoring ordinary lexical ambiguity, mixing of metaphor with other figurative phenomena, and other complications.

Section 1 that the source and target domains there overlap – they both involve people, mental states and (often) natural language communication.

But quite apart from overlap, there is no convincing way of dividing knowledge up into domains anyway. The often-implicit claims about domains being different in the metaphor literature are often self-serving, context-sensitive and without independent justification. For instance, in a metaphorical portrayal of a person as a fox, the presumption (by someone who wanted to distance source domains from target domains) would have to be that foxes and people are in different domains. But, although we certainly do qualitatively distinguish sometimes between animals and people in common sense, we would lump them together in the case of some types of metaphor. For instance, in the IDEAS AS ANIMATE BEINGS metaphorical view, the animate beings could be either people or animals, and sometimes there may be no clue as to which or importance in making the decision (consider, e.g., “The idea was lurking in the back of his mind”). In any case, the domain of ideas presumably involves people, because people are the (main or only) beings that have ideas, so the source and target domains would overlap.

Moreover, the boundaries of a domain are extremely difficult to determine. What should be in the domain of people, for instance? Should electrons be, because they are parts of people? Should for that matter foxes be, because people hunt foxes? Questions like these have led the design of ATT-Meta to shy away from any requirement to label particular rules, predicates, etc. as belonging to particular domains. Therefore, there is no way for the system to control, on the basis of domains, what rules, etc., can be used in a particular pretence cocoon.

However, the matter goes deeper than this. It is incorrect to think of a pretence cocoon as being tied in principle to the source domain of the metaphor at hand: it is just that in practice, because of the action of source-to-target rules associated with metaphorical views, subqueries that are in source-domain terms are placed in the cocoon, and in practice most, but not necessarily all, of the resultant backwards reasoning might be *theoretically* labelled, by an observer who ascribes to a particular set of domains, as being within the source domain.

4 What ATT-Meta Does Overall

ATT-Meta itself has no knowledge of any specific metaphorical view or any particular domain. It is merely a reasoning engine. It does, however, include some built-in rules about pretence, beliefs, qualitative degrees, and some other matters. The user supplies, as data for ATT-Meta, whatever target-domain knowledge, source-domain knowledge, metaphorical mapping relationships and ancillary assumptions the user wants to try out for handling examples. These pieces of information are expressed as rules (with fact-rules as a special case). The rules are applied by ATT-Meta. In the illustrative syntax used in this article, a simple rule about birds might be

```
IF bird(X) AND alive(X) THEN {presumed} can-fly(X).
```

The symbol `presumed` is a qualitative certainty qualifier (see below), and can also be read as “by default.” Rules with this qualifier are intuitively similar to “normal” rules in Default Logic (Reiter 1980). If-then rules with a special, essentially null IF part count as fact-rules. We will write a fact-rule without any IF or THEN part, as in

`{certain} is-person(Anne) .`

The user supplies fact-rules that are intended to couch the logical forms of the direct meanings of one utterance or a small set of utterances that form a (real or imaginary) chunk of discourse. Utterances may or may not be metaphorical. In the metaphorical case, the fact-rule contains a special feature (see Barnden & Lee 2001) that causes the factual hypothesis arising from it is placed in the pretence cocoon, not the top reasoning context. Such factual hypotheses can be seen at the top of the cocoon box in Figure 1.

The user supplies a top reasoning query that ultimately drives all the reasoning. For example, the query could be “*Anne is consciously aware to degree at least D of the idea that Kyle is unfaithful?*” In reality, we suppose that surrounding discourse has the effect of posing such queries (or queries for which they are indirect subqueries). For instance, it may be that surrounding discourse makes Anne’s conscious mental state about Kyle be of interest.

When the top query does not contain a variable (such as the D just above), ATT-Meta tries to find evidence for and against the top query, using the user-supplied knowledge and the user-supplied logical forms of utterances. In cases of conflict, a conflict-resolution mechanism attempts to favour one side or the other. When the top query contains variables, ATT-Meta tries to find values for the variables such that the thereby instantiated versions of the query are supported.

5 Hypotheses and Rule Application

The information dynamically manipulated by ATT-Meta consists of *hypotheses*. Hypotheses are reasoning queries or propositions created by rule applications. A query is either the top query supplied by the user or a subquery created by ATT-Meta in the course of backwards reasoning in relation to the top query. One simple way a query can be supported is if it matches a user-supplied fact-rule. The query hypothesis is then called a *factual hypothesis* or just a *fact*. Queries can contain variables, and these are regarded implicitly as existentially quantified. Thus, intuitively, the query `can-fly(X)` is asking whether there is at least one entity that can fly. The system tries to find all possible values for X for which `can-fly(X)` is supported.

At any time, any particular hypothesis H is tagged with a qualitative certainty level, one of `certain`, `presumed`, `suggested`, `possible` or `certainly-not`. The last one 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 H, but the evidence is not (yet) strong enough to enable H to be a working assumption. When a hypothesis is created (as a query), it is immediately given a certainty value of `possible`. Of course, reasoning may sooner or later downgrade it to `certainly-not` (i.e., impossible).

When a rule is applied, the certainty it contributes to its result is the minimum of its own certainty qualifier and the certainty levels assigned to the hypotheses picked up by the condition part. Multiple rule applications can support a given hypothesis. In this case, the maximum of the certainty values contributed by the different applications is used. When there is evidence to level `presumed` for both a query and its complement

(complement hypotheses being negations of each other), then at an appropriate point the conflict-resolution mechanism steps in and tries to adjudicate the relative evidence strength.

Hypotheses are annotated as to which reasoning contexts they exist in. In this article, we need consider only two contexts: a single pretence cocoon and the reasoning environment outside the pretence cocoon, which we call the “top” context. The top context is the system’s own reasoning context, in which the hypotheses are about reality as the system sees it. A hypothesis can be replicated in different contexts, but if so the different copies are handled entirely separately, as if they were unrelated hypotheses. This is because even if a hypothesis is supported in two different contexts, the lines of reasoning supporting it can differ radically, and the level of certainty can differ.

Except in the case of metaphor-related rules (including source-to-target mapping rules), when a rule is applied to a query in a specific reasoning context, the subqueries emanating from its condition part are set up in the same reasoning context except when special rule features dictate otherwise (Barnden & Lee 2001). All rules are available for use in any context, but rules can contain special conditions that can prevent the application of the rule if the context is inappropriate (Barnden & Lee 2001).

6 Basic Conflict Resolution

Here we describe the basic conflict-resolution provisions. There are additional provisions concerning pretence cocoons that are in Section 7 below, and they are our main target in this paper, given its orientation towards contexts. The provisions in the present section apply to rules of any sort, in any context. However, for simplicity we will explicitly talk about them being applied only to adjudicate conflicts between reasoning chains arising in the top reasoning context. Various complications in basic conflict resolution will not be addressed in this paper:

- The question of exactly when conflict resolution should happen, in cases where the derivation network for hypotheses contains cycles. There is a complex algorithm for postponing conflict resolution until all hypotheses around a cycle can be treated together and fairly.
- The fact that `presumed` hypotheses can be non-downgradable because they depend via `certain` rules on `presumed` hypotheses whose certainty level has already been finalized.
- Back-propagation of downgrading through applications of `certain` rules, to help maintain consistency of certainty levels.
- Dealing with the point that facts can sometimes also be derivable from other facts. Normally, in such cases the derivable facts should be treated as non-facts for the purpose of conflict-resolution.

Recall that conflict resolution steps in when both a hypothesis `H` and its complement are supported to level `presumed`. The final result of the resolution attempt will be that one or both hypotheses will be downgraded to `suggested`. They are both downgraded if neither can be determined to be the winner. On the other hand, when one is the winner, it stays at `presumed`, and the `suggested` level that is given to the other one records

the fact that there *was* some strong evidence for it. Fundamentally, conflict resolution *specificity-based* and is to that extent traditional (cf. exception handling in semantic networks, and see, e.g., Loui 1987). That is, the system tries to see whether one of the hypotheses relies on a more specific situation than the other one does. This comparison is based on the sets of facts employed by the arguments supporting the hypotheses, combined with some consideration of the precise way the hypotheses depend on those facts. We now describe the core of the specificity comparison.

6.1 Core of the Fact-Based Specificity Comparison

The main intuition of the fact-based specificity comparison is that if one of the conflicting hypothesis relies on more facts, in a set-theoretic rather than numerical sense, then it is more specifically supported. For instance, if the hypothesis that Ralph is not tidy follows only from the fact that he is a student, whereas the hypothesis that he *is* tidy follows from the combination of the facts that he is a student and is elderly, then the latter hypothesis wins. To be more precise, suppose ATT-Meta contains the following rules, where *at-least-m-a* means *at-least-middle-aged*:

```
IF student(X) THEN {presumed} NOT(tidy(X)).
IF student(X) AND at-least-m-a(X) THEN {presumed} tidy(X).
IF elderly(X) THEN {certain} at-least-m-a(X).

{certain} student(ralph).      {certain} elderly(ralph).
```

It is clear that there is one argument for the hypothesis RNT that Ralph is not tidy, and that this argument relies on the set of facts { *student(ralph)* }. There is also one argument for the hypothesis RT that Ralph is tidy, and it relies on the factset { *student(ralph)*, *elderly(ralph)* }. Both facts in that set are necessary for the argument. Now, in other examples, there could be alternative arguments for either hypothesis, giving alternative factsets. ATT-Meta therefore calculates the *factset-set* (set of factsets) for each of the conflicting hypotheses. In our example, because there is just one argument for each hypothesis, we have just:

```
factset-set for RNT:      { {student(ralph)} }
factset-set for RT:      { {student(ralph), elderly(ralph)} }.
```

This will make RT win, because of the factset-set comparison method detailed below. Here, it is tantamount to noting that the single factset for RNT is a subset of the single factset for RT, but the factset for RT is not a subset of the factset for RNT. The single argument for RNT relies on a situation that is at least as unspecific (actually, strictly less specific) than the single argument for RT; but the reverse is not true.

When a hypothesis can be supported by more than one rule application, these applications are parts of alternative arguments for the hypothesis. They are therefore also parts of alternative arguments for other hypotheses depending on that hypothesis. When there are alternative arguments for a hypothesis, where the alternatives rely on different facts, it generally has more than one supporting factset. We see a simple illustration of this in our Ralph example augmented with the following rules:

```
IF musician(X) THEN {presumed} NOT(tidy(X)).
{certain} musician(ralph).
```

The two factset-sets are now

factset-set for RNT: $\{ \{ \text{student}(\text{ralph}) \}, \{ \text{musician}(\text{ralph}) \} \}$
factset-set for RT: $\{ \{ \text{student}(\text{ralph}), \text{elderly}(\text{ralph}) \} \}$.

In this case, there is no winner because there is no clear sense in which one of the hypotheses has more specific support overall than the other. The reason RT won before was that its single factset was not a subset of any factset for RNT. This is still true, but the converse is now also true: the new factset for RNT is not a subset of any factset for RT. Intuitively, there is no support for RT that is at least as specific as the new argument for RNT. It seems reasonable to say that for there to be a losing hypothesis, *every* argument for it must rely only on facts that are also used by *some* argument for the winner. To respect this principle, the actual specificity comparison on factset-sets is as follows.

Definition: Let the two factset-sets for a hypothesis and its complement be FSS1 and FSS2. Then FSS2 is deemed to be more specific than FSS1 if and only if

- for every member FS of FSS1, FS a subset of some member of FSS2
- the reverse is not true.

An abstract example of FSS2 being more specific than FSS1 is

FSS1 = $\{ \{ a,b \}, \{ b,c,d \}, \{ a,b,e,f \}, \{ k \} \}$
FSS2 = $\{ \{ a,b,e \}, \{ a,b,e,f,g \}, \{ b,c,d \}, \{ k,l \}, \{ m \} \}$

Even though one member of FSS2, namely $\{ a,b,e \}$, is a (proper) subset of a member of FSS1, the conditions for FSS2 being more specific still hold, because at least one member, in fact three members, of FSS2 are not subsets of any member of FSS1. Notice also that FSS1 and FSS2 can intersect: $\{ b,c,d \}$ is a member of both. In fact, an extreme case of this, which can arise in practice, is that FSS2 is more specific than FSS1 even when FSS1 is just a proper subset of FSS2, i.e. they contain the same factsets except that FSS2 has at least one more. An example:

FSS1 = $\{ \{ a,b \}, \{ b,c,d \} \}$
FSS2 = $\{ \{ a,b \}, \{ b,c,d \}, \{ a,c \} \}$

even though the unshared fact-set $\{ a,c \}$ in FSS2 happens not to introduce any facts that do not appear in the shared factsets. But it is still reasonable for FSS2 to be deemed more specific, because it uses facts a and c in an extra way that is not paralleled by any factset in FSS1.

The factset-set comparison provision is quite a powerful technique, but there is a widespread type of phenomenon it fails on. Consider the question of whether penguins can swim underwater. If we have

```
IF bird(X) THEN {presumed} NOT(swim-uw(X)).
IF penguin(X) THEN {presumed} swim-uw(X).
IF penguin(X) THEN {certain} bird(X).
IF says(mary, penguin(X)) THEN {presumed} penguin(X).
{certain} says(mary, penguin(peter)).
```

then both conflicting hypotheses (that Peter can swim underwater, and the negation) have just one factset, $\{ \{ p \} \}$, where p is the fact that Peter is a penguin. The way ATT-Meta handles this situation is given in the next subsection.

where m is `says(mary, penguin(peter))`, p is `penguin(peter)` and b is `bird(peter)`. Although $\{(m, p)\}$ is a shared prefix, it is not independent for hypothesis `NOT(swim-uw(peter))` because the only supporting path for that hypothesis that contains that prefix also contains the more extensive shared prefix $\{(m, p, b)\}$.

The next step is find the intersection of the two sets of independent shared prefixes, namely $\{(m, p, b)\}$ in our example. For each prefix in the intersection, the last link in the prefix is removed from the derivation network.

The technique is quite robust, in the sense that extra hypotheses can interrupt the links on the original derivation network without changing the specificity comparison. In addition, arguments for `NOT(swim-uw(peter))` that do *not* pass through the `bird(peter)` hypothesis are not affected by the particular manipulations above, as is intuitively correct, as it allows `NOT(swim-uw(peter))` then possibly to win. The technique also copes easily with iterated defeat, as arises, for instance, if sick penguins cannot normally swim underwater (even though they are penguins), but pregnant sick penguins can (even though they are sick penguins). We do not have space to demonstrate this here.

7 Matters of Pretence: Rules, Fact Importation and Pretence Bias

We noted in Section 3 that any non-fact rule possessed by the system can be considered for application inside the pretence cocoon. But, because of the unreal situations that can arise in pretence, rules that have strength `certain` are downgraded to strength `presumed` when used inside the pretence cocoon, to allow more intuitive conflict-resolution. The reason for not allowing fact-rules to be used in the pretence cocoon, unless they are specially written to specify this, is that typically such rules describe the specific circumstances holding in a real situation or holding within the pretence. However, because of the possible overlap of source and target domains in a metaphor, and because of the blending that happens in any case by virtue of the initialization of the pretence (by a fact saying, for instance, that Anne's mind has far reaches), ATT-Meta provides for *importation* of facts holding in the top context into the pretence cocoon. Note that this is importation of factual hypotheses, not a use of fact-rules within the pretence, so that what happens is that two copies of the fact arise, one in the cocoon and one outside, and the former depends on the latter via an *importation link*. One is shown in Figure 1.

Also, importation can support the inner copy only to level `presumed`, much for the same reason as that non-fact rule certainties are capped at `presumed`. But this capping of fact certainty is not enough to avoid unintuitive conflict resolutions, and imported facts need to be given lower priority in the conflict-resolution process inside the cocoon than facts that are inside the cocoon, or than facts that are outside the cocoon and support within-cocoon hypotheses via the action of target-to-source rules (such as ancillary assumption rules). Both these types of facts should be more powerful in providing specificity, because they are to do with special features of the metaphor at hand, whereas importation is a completely general act insensitive to the particular metaphor at hand. An illustration is given by the Anne/Kyle example, as follows.

There is an argument inside the cocoon that Anne's mind is presumably a physical space. The fact this argument rests on is that Anne's mind has far reaches, as illustrated in Figure 1. This fact arises directly in the cocoon by a fact-rule—it is not the result of an importation. But in the Anne/Kyle example (and of course others) we provide rules saying that minds are certainly *not* physical spaces, and that for, any person X , $\text{mind-of}(X)$ is a mind; and in the example we have the fact-rule that Anne is a person. Thus, by virtue of the importation of that fact that Anne is a person into the cocoon, an argument arises in the cocoon that Anne's mind is presumably not a physical space. (This is not shown in the Figure.) We therefore have conflict. Because the importation would have been available no matter what the metaphorical utterance had been, what the metaphorical views in play had been, or what specific situation had been under consideration, we wish to make the importation-based argument look less specific to the conflict-resolution than the non-importation-based argument for the hypothesis that Anne's mind is a physical space.

In more complex situations, it can be that in a given factset supporting a hypothesis there are both imported facts and non-imported ones. Different factsets for a hypothesis can differ on whether they contain imported facts or not. Both hypotheses in the conflict can have factsets containing imported facts. And, an imported fact might have independent, non-importation-based support within the cocoon. The general handling of anti-importation bias takes these complexities into account by first ignoring the contribution of imported facts, and only using them in the specificity comparison if the conflicting hypotheses are otherwise equal in specificity (i.e., have equal factset-sets). One desirable effect of this is that a hypothesis that is partly, but only partly, dependent on importation will win against its complement if the latter is entirely dependent on importation. Also, importation cannot upset a winning situation that arises from the special circumstances arising within the pretence.

In more detail, ATT-Meta's factset-set comparison happens in (possibly) two phases. In the first phase, importation is ignored. That is, the derivation network is slimmed down by removing all links that represent importation acts; any hypotheses therefore deprived of all support are also removed, and so on. A rule application is only removed from the derivation network when all its condition nodes are removed. After this slimming process, the factset-sets are determined and compared. If one hypothesis wins, the conflict resolution is complete and there is no second phase.

Otherwise, if the two factset-sets are unequal, so that the two hypotheses have *incomparable* specificity, the conflict resolution attempt is again complete and there is no second phase, but this time there is failure to find a winner. However, if the two factset-sets are the same, so that the two hypotheses are *equally* specific, the second phase is executed, and a winner then possibly found. In the second phase, the original derivation network from before the slimming process is reinstated, and factset-sets again calculated and compared.

In the Anne/Kyle example, the hypothesis that Anne's mind is not a physical space only has one fact in its pre-slimming factset, and this fact is the imported one that Anne is a person. Therefore, in the first phase this factset-set becomes empty. On the other hand, the hypothesis that Anne's mind is a physical space relies on the non-imported

fact that it has far reaches, and the slimming therefore does not affect this factset-set. The latter hypothesis therefore wins in the first phase and there is no second phase.

The two-phase process as described here is actually a special case of a more general process that caters for nesting of cocoons within each other (Lee & Barnden 2001) by given preference for deeper importations over shallower ones, but this is beyond the scope of the present paper.

8 Conclusion

We have argued that ATT-Meta can robustly handle conflict and uncertainty in both the reality and pretence contexts, and in particular can cope with conflicts in the pretence cocoon between source-domain reasoning based on special pretence circumstances and conclusions depending partly on non-pretence information. ATT-Meta can handle other types of uncertainty and conflict (such as conflict between intrinsic target information and the results of metaphorical reasoning), but this will be presented in detail elsewhere.

References

- Barnden, J.A. (1998). Combining uncertain belief reasoning and uncertain metaphor-based reasoning. In *Procs. Twentieth Annual Meeting of the Cognitive Science Society*, pp.114–119. Mahwah, N.J.: Lawrence Erlbaum Associates.
- 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.
- Barnden, J.A. & Lee, M.G. (1999). An implemented context system that combines belief reasoning, metaphor-based reasoning and uncertainty handling. In P. Bouquet, P. Brezillon & L. Serafini (Eds), *Lecture Notes in Artificial Intelligence, 1688*, pp.28–41, Springer. (CONTEXT'99.)
- Barnden, J.A. & Lee, M.G. (2001). Understanding usages of conceptual metaphors: An approach and artificial intelligence system. Technical Report CSRP-01-05, School of Computer Science, The University of Birmingham, U.K.
- Creary, L. G. (1979). Propositional attitudes: Fregean representation and simulative reasoning. *Procs. 6th. Int. Joint Conf. on Artificial Intelligence (Tokyo)*, pp.176–181. Los Altos, CA: Morgan Kaufmann.
- Davies, M. & Stone, T. (Eds) (1995). *Mental simulation: evaluations and applications*. Oxford, U.K.: Blackwell.
- Kittay, E.F. (1989). *Metaphor: its cognitive force and linguistic structure*. (Paperback ed.) Oxford, U.K.: Clarendon Press.
- Lakoff, G. (1993). The contemporary theory of metaphor. In A. Ortony (Ed.), *Metaphor and Thought*, 2nd ed., pp.202–251. Cambridge, U.K.: Cambridge University Press.
- Lee, M.G. & Barnden, J.A. (2001). Reasoning about mixed metaphors with an implemented AI system. *Metaphor and Symbol*, 16 (1&2), pp.29–42.
- Loui, R.P. (1987). Defeat among arguments: a system of defeasible inference. *Computational Intelligence*, 3, pp.100–106.
- Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13, pp.81–132.

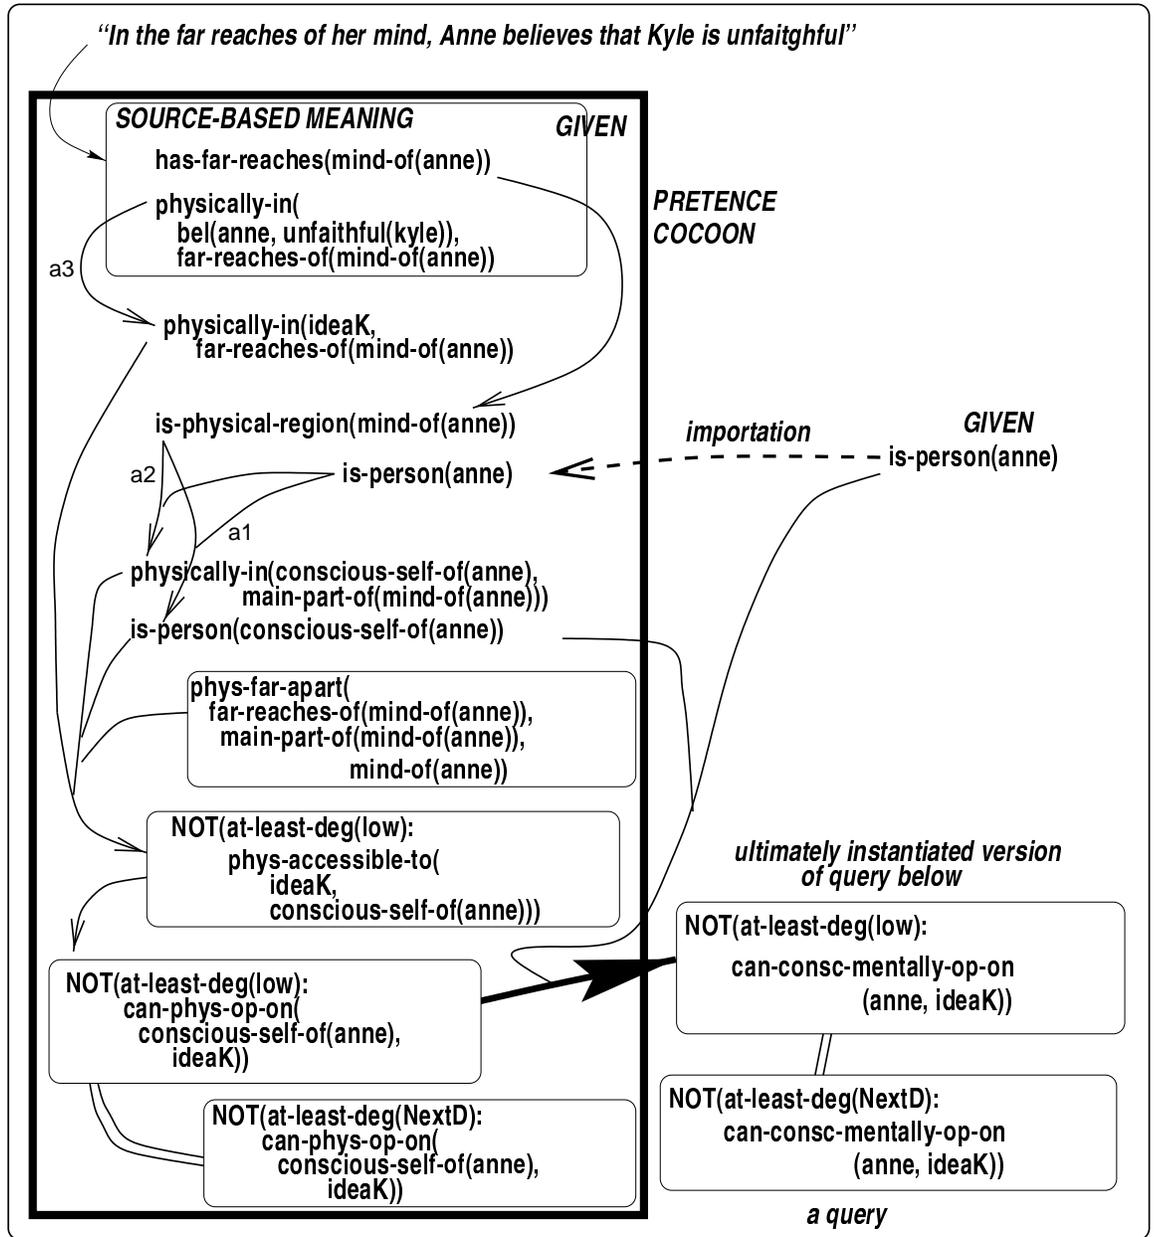


Fig. 1. Showing the main features of ATT-Meta’s metaphor-based reasoning in a particular application that was performed on the Anne/Kyle example. The logical notation is an illustrative, simplified form of that used in the system. The symbol *ideaK* is an abbreviation for the-idea-that(unfaithful(kyle)). The third argument of the *phys-far-apart* hypothesis specifies that the far-apartness (of the far reaches and the main part) is relative to the scale of Anne’s mind as a whole. The heavily outlined box depicts the pretence cocoon. Facts are marked as GIVEN; both hypotheses in the small box at top left are facts. All facts in the example are certain. All other hypotheses shown attain a certainty qualifier of presumed. The fat arrow depicts the application of source-to-target metaphorical mapping rule. Applications of ancillary-assumption rules are numbered a1–3. Double lines connect general queries with their instantiations. The depiction of some rule applications is simplified (i.e., links omitted), and many hypotheses and reasoning steps are omitted, including the within-pretence hypothesis that *ideaK* is a physical object. This hypothesis is supported within the pretence by the hypothesis that *ideaK* is in a physical region. The query at bottom right is a subquery arising from the actual top query used, which was “Can Anne consciously operate mentally on IdeaK to exactly degree D?”