Application of the ATT-Meta Metaphor-Understanding System to an Example of the Metaphorical View of MIND PARTS AS PERSONS

Technical Report CSRP-01-09

John A. Barnden & 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

The main aim of this report is to provide a measure of objective evaluation of ATT-Meta, a system we have implemented for conducting some of the reasoning needed in the understanding of metaphorical utterances. The report thereby also provides some evaluation of the overall theoretical approach informing ATT-Meta. The approach and system are not described in detail here, but rather in an accompanying technical report (CSRP-01-05). However, the present report is self-contained. Its evaluation of the ATT-Meta system consists of showing how the system extracts important target-domain information from a particular metaphor example that was found in real discourse. This example rests on the metaphorical view of MIND PARTS AS PERSONS, in which a person or a person's mind is viewed as having parts that are themselves persons, or as otherwise containing persons, where these persons have their own mental states. The view is extensively used in mundane conversation and text. The report shows that the system successfully produces major pieces of information conveyed by the example utterance. It does so on the basis of a very limited set of mapping relationships included in ATT-Meta as its long-term knowledge of the metaphorical view. The implementation of the example also helps to vindicate the Map-Extension Minimization stance within the ATT-Meta approach. This stance is that metaphorical mappings should, by default, not be extended to cope with utterance aspects that those mappings do not deal with. Rather, within-source-domain reasoning should be used to link those aspects to source-domain aspects that the mappings do already deal with.

1 Introduction

The main aim of this report is to provide a measure of objective evaluation of ATT-Meta, a system we have implemented for conducting some of the reasoning needed in the understanding of metaphorical utterances. The report thereby also provides some evaluation of the overall theoretical approach informing and partially implemented in the ATT-Meta system. The evaluation of the ATT-Meta system in the present report consists of showing how the system extracts important target-domain information from a particular metaphor example that we found in real-discourse. This example rests on the metaphorical view of MIND PARTS AS PERSONS (OR OTHER ANIMATE BEINGS), in which a person or a person's mind is viewed as having parts that are themselves persons, or as otherwise containing persons; these subpersons have their own mental states and are often portrayed as communicating with each other in natural language. The subpersons' mental states/processes are mapped in a particular way to mental states/processes of the whole, real agent, and intercommunication between the subpersons is mapped to a mental process of the whole agent.

The approach and system will not be described in detail here—the reader is referred to other papers, especially another technical report, Barnden & Lee (2001). (See also Barnden *et al.*, 1996; Barnden, 1998a,b; Barnden & Lee, 1999). The present report is intended as an adjunct to Barnden & Lee (2001), and although the present one is self-contained it is ideally to be read after absorbing the less technical portions of the other report.

The metaphorical utterances of most interest in the ATT-Meta project are ones we call *map-transcending*: they rest on metaphorical views (conceptual metaphors in the terminology of Lakoff, 1993) that are familiar to the understander, but transcend them by using source-domain concepts not handled by the between-domain mappings the metaphorical views involve. A possible example is "McEnroe starved Connors to death" when used as a description of a tennis match. A metaphorical view of SPORT AS COMBAT might be familiar to the understander, and that view might map combat-domain concepts such as killing to target-domain concepts such as defeating. However, the existing mapping might not contain any mapping relationship that maps the source-domain concept of starving someone. (We regard a mapping to be a set of mapping relationships, each of which maps an aspect of the source domain to an aspect of the target domain or vice versa.)

For map-transcending utterances based on familiar metaphorical views, the ATT-Meta approach advocates possibly-extensive inferencing in the terms of the source domain of the metaphorical view, while avoiding as far as possible the extension of the mappings to deal with the map-transcending source-domain aspects of the utterance. This is the Map-Extension Minimization default. Map-extension is regarded as an exceptional action undertaken only when special needs of discourse understanding force it to take place.

The source-domain inferencing is done in an attempt to derive, from the source-domain meaning of the utterance, information that *can* be converted by the existing source-to-target mapping into information about the target domain. The source-domain meaning is the meaning the utterance has when taken at face value as being about the source domain itself. The source-domain inferencing is done within a special, protected computational environment we call a *pretence cocoon.* See Figure 1. The cocoon can heuristically be thought of as the belief space of a hypothetical agent who believes that the source-domain meaning of the utterance is true. In the above starving example, the source-domain meaning is that McEnroe starved Connors to death in the real, biological sense. The inferencing within the cocoon can then establish that McEnroe killed Connors. The existing mapping can then be applied to derive, in the target domain and outside the cocoon, the hypothesis that McEnroe defeated Connors in the match. Other inferences could be made within the cocoon on the basis of the detailed form of killing that starving someone constitutes, and some resulting inferences might be mapped over to the target domain by additional existing mapping relationships.

There is a great deal of understander-relativity in the ATT-Meta approach to metaphor understanding. For one thing, the question of whether a metaphorical view is familiar or not is an understander-relative issue. However, we assume that the particular views mentioned in this report are familiar in some form to ordinary English speakers (which does not mean that speakers are consciously aware of them). The assumption is based partly on the high frequency with which we have observed the views being used in mundane discourse, partly on intuition, and partly on the observations of other metaphor theorists (see particularly Lakoff 1996 for metaphorical views strongly related to MIND PARTS AS PERSONS). See also our metaphor-example databank (Barnden, n.d.) for real-discourse examples of MIND PARTS AS PERSONS and other metaphorical views of mind. See Barnden (1997) for general discussion of MIND PARTS AS PERSONS.

There is further understander-relativity in the question of what particular mapping relationships between source and target are included in viewing something A as something B. For example, different understanders may differ in detail on what the metaphorical view of MIND PARTS AS PERSONS amounts to. We have tried to include in any given metaphorical view that we have addressed in the ATT-Meta approach only those source/target mapping relationships that seem to be needed in many examples we have encountered and that seem to us intuitively acceptable.

We also assume that stock English metaphorical phraseology (i.e., phraseology that is frequently occurring and quite fixed, such as "see the problem") uses metaphorical views that are familiar to typical English understanders, and that (in general) the source-domain aspects used in the phraseology are mapped by the mappings in the views. Hence, we assume that stock phraseology does not count as map-transcending.

Additionally, we will only consider examples in which the metaphoricity is plausibly not *sidelined* for an ordinary English speaker. For example, an understander might have a sports-domain lexicon sense for "kill," so that he/she/it could use that sense in understanding the sentence "McEnroe killed Connors" in the tennis context. This would circumvent the need to proceed via a sense of kill in the biological combat domain and a mapping from that domain to the domain of sport. Clearly, whether the metaphoricity of an utterance is sidelined or not is again an understander-relative matter.[1]

For a particular metaphorical view, the class of utterances that use stock phraseology associated with the view and the class for which the metaphoricity is sidelined can naturally be expected to overlap to a high degree. However, there is no logical necessity for the metaphoricity of stock phraseology to be sidelined or for utterances having entirely sidelined metaphoricity to use only stock phraseology. Hence, it is useful to preserve the distinction between the two notions.

The evaluation of the ATT-Meta system in this report consists of showing how the approach could extract important target-domain information from a somewhat simplified version of a plausibly-non-sidelined, plausibly-map-transcending metaphorical utterance that we found in real, mundane discourse. The example is in our databank (Barnden, n.d.), where, in particular, other examples from the same pulp-romantic story that the chosen example was taken from may be found.

In our work on the ATT-Meta approach we have been mostly concerned with examples that are from mundane discourse, or that at least plausible candidates for being found in mundane discourse. By mundane discourse we mean discourse in conversations, popular magazine articles, popular novels, news articles, popular science texts and other factual texts for wide consumption by the general public. However, the approach is also applicable to non-mundane discourse.

The ATT-Meta approach is not as it stands a specific, complete model of metaphor understanding. It is parametrized by whatever views are familiar to a particular understander to whom/which we are imputing

[1] Observe that we can still regard "McEnroe killed Connors" as being metaphorical for the understander as long as the biological-combat-domain route is in principle *available* to the understander.

the approach, what mapping relationships are included in each familiar view, and what direct, lexicon senses the understander has for words. Thus, it is important to understand that the purpose of this report is not to *establish* that any particular metaphorical views are familiar to typical understanders, that any particular mapping relationships appear in particular views, that the metaphoricity in any particular example is not sidelined for typical understanders, or other such matters. The purpose is to show that, given particular decisions about what is familiar, what is mapped what is not sidelined, etc., the ATT-Meta approach can succeed in drawing useful information about the target domain from a metaphorical utterance; and also to show that very little has to be included in the way of mapping information in order to get great power in drawing out such information. The less sidelining we assume and the fewer mapping relationships we include in a metaphorical view, the more the power of the approach is demonstrated.

We believe the discussion of the particular example below covers the main "informational contributions" that the chosen example makes concerning its target domain. We doubt that the suggested informational contributions are controversial.

The plan of the rest of the report is as follows. Section 2 summarizes an important additional feature of the treatment of source/target mappings in the ATT-Meta approach. Section 3 mentions some key features of the ATT-Meta system. Section 4 introduces the original form of the metaphor example with which this report is concerned. In this section we discuss how the general approach, rather than the system itself, could handle the example. Section 5 shows how the system applies to a simplified form of the example. The system is not able to achieve all that the approach is intended to, but captures some important aspects of the example. Section 6 concludes.

Other articles provide additional evaluation. Lee and Barnden (2001b) evaluate the implemented system and the general approach by applying the system to real-discourse utterances based on most of the metaphorical views about mental state/processes that are listed in the Master Metaphor List (Lakoff, 1994). See Barnden & Lee (2001) for a detailed description of how the system works on another real-discourse example from the databank. The examples in these various papers concentrate on metaphor of mind, but the approach is in no way limited to such metaphor. Barnden (2001) evaluates the approach (but not the system) by applying it to selected, mostly non-mental, examples from Goatly's (1997) comprehensive, detailed survey of metaphorical language.

2 View-Neutral Mapping Adjuncts

Barnden & Lee (2001) specify various *view-neutral mapping adjuncts* (VNMAs) that form part of the general ATT-Meta approach. These provide additional mapping relationships that apply, by default, whatever particular metaphorical views are at hand. They are parasitic on other things that are mapped, in the sense that they are all of the form: if such-and-such things are mapped from the pretence into reality space, then also such-and-such things are mapped (by default).

The VNMAs do not have as yet a full and satisfactory implementation in the ATT-Meta system. Some are included in a limited, provisional way, and we discuss this point in the Conclusion section. Since in our general discussion of the chosen example we will be appealing to VNMAs, we summarize them here.

One of the simplest VNMAs is the *Negation* VNMA, which states that if a within-pretence proposition P is mapped into reality space to become a proposition R there, then the negation of P is also mapped, and is mapped to the negation of R. For example, if the proposition that McEnroe killed Connors is mapped to the proposition that he defeated Connors, then, by virtue of the VNMA, the view also maps the proposition that McEnroe did **not** kill Connors to the proposition that he did **not** defeat Connors. Relatedly, the *Qualitative Degree* VNMA maps identically the qualitative degrees to which mapped properties and relationships hold.

The *Causation/Ability* VNMA identically maps properties and relationships such as causation, enablement, disenablement, ability, assistance, hindrance and easiness. For example, consider the within-pretence proposition P that McEnroe's combat actions **enabled** him to *kill* Connors **easily**. If the metaphorical view maps combat actions to tennis actions and killing to defeating, then P is mapped to the proposition in reality-space that McEnroe's tennis actions **enabled** him to *defeat* Connors **easily.**

The *Mental/Emotional States* VNMA maps the mental/emotional states of within-pretence agents over to mental/emotional states of corresponding reality-space agents, if any. It does so as long as the propositional content (if any) is itself mapped. The mental attitude (e.g. belief) or the quality of the emotion (e.g. anger) is preserved, but the propositional content (if any) may be modified by other mapping relationships. For example, if an entity E is an agent both inside and outside the pretence in the McEnroe/Connors example, and in the pretence E **believes** that Connors *killed* McEnroe, then in reality E **believes** that Connors *defeated* McEnroe.

Somewhat similarly, the *Value-Judgment* VNMA identically maps value-judgments (whether by the understander, other real agents, or within-pretence agents) about mapped states of affairs, issues, etc. within the pretence to value-judgments about the corresponding states of affairs, issues, etc. in reality. Value judgements include judgments about levels of goodness and importance.

As we do not, in general, regard any of the "subpersons" in the MIND PARTS AS PERSONS view as mapping to the real person, the Mental/Emotional States VNMA does not just by itself provide us with a mapping relationship that deals with the mental states of subpersons. However, some utterances based on the view do clearly involve an identification of one of the subpersons as representing at least the main thoughts of the real person, and this subperson could be termed the conscious self of the person. So the VNMA does deal with the mental states of that subperson.

Further VNMAs and more refined descriptions are given in Barnden & Lee (2001). Our VNMAs are inspired in part by, but are different from, postulates about mapping put forward by other authors, as detailed in Barnden & Lee (2001).

3 The ATT-Meta System

This paper will merely sketch some aspects of the ATT-Meta system, to the extent needed for the current paper, but further details can be found in (Barnden et al., 1996; Barnden, 1998a,b; Barnden & Lee, 1999; Barnden & Lee, 2001; Lee & Barnden, 2001a).

In ATT-Meta, reasoning is done by the use of back-chaining rules of inference which allow differing qualitative degrees of certainty. Reasoning in ATT-Meta is entirely query-directed (goal-directed), rather than proceeding forwards from data. The user supplies some query which would in reality be generated as a result of partial processing of surrounding discourse or other context. This is explained further in Barnden & Lee (2001).

ATT-Meta is capable of both simulative reasoning about beliefs and metaphorical reasoning. Nested reasoning spaces ("cocoons") are allowed to facilitate simulation of other agents and to facilitate metaphorical reasoning. Two types of cocoons are maintained: simulation-pretence cocoons (also called belief cocoons) and metaphor-pretence cocoons. Simulation-pretence cocoons are used to model the beliefs of other agents. Metaphor-pretence cocoons are a special type of simulation-pretence cocoon where the agent modelled is a hypothetical agent who is assumed to believe that the metaphorical utterance is literally true. For the remainder of this paper we will only be concerned with metaphor-pretence cocoons, although belief cocoons also arise in a minor way in the implemented example.

Knowledge of different domains is encoded in sets of rules which apply to a particular domain. Since a metaphorical view involves mapping relationships between one domain (the source) and another (the target), ATT-Meta uses "conversion rules" which explicitly map propositions from one domain to another. ATT-Meta has a small set of conversion rules for each metaphorical view it is given knowledge about, as well as holding knowledge about the source domain of each such view. (ATT-Meta's conversion rules can go from target to source as well as from source to target, for reasons given by Barnden, in press.)

In the ATT-Meta approach, a metaphorical view can be defined by constructing a small set of conversion rules and a set of rules to represent the source domain (and sometimes a small set of "ancillary assumption" rules that are somewhat similar in form to conversion rules—see Barnden & Lee, 2001). Understanding of a metaphorical utterance proceeds by creating a metaphor-pretence cocoon where the utterance is taken as literally true, then finding implications of this within the cocoon, largely by means of source-domain reasoning, and then mapping some implications to the target domain via conversion rules. The within-cocoon derivation process can involve a substantial amount of reasoning.

ATT-Meta is merely a reasoning system, with no natural language front-end as yet. Thus, the user supplies hand-coded logical expressions that are intended to couch the source-domain meaning of the imagined input utterance, and other logical expressions couching additional information that is assumed to arise from surrounding discourse. In addition, the user supplies the domain rules, conversion rules and ancillary-assumption rules—they are not built into the ATT-Meta reasoning engine itself. Naturally, a given set of conversion rules is thought of as being stable over the "life" of the system, and to be able to handle many examples.

In ATT-Meta, hypotheses are tagged with qualitative certainty levels, which are largely suppressed in this report for the sake of simplicity of presentation. The inferred conclusions presented in examples below all achieve the certainty level called `presumed`. The available levels are as follows, with their meanings when attached to a proposition H:

`certain:` H is established with complete certainty.

`presumed:` H is a default: i.e., it is taken as a working assumption, pending further evidence.

`suggested:` there is evidence for H, but the evidence is not (yet) strong enough to enable H to be a working assumption.

`possible:` the negation of H is not `certain` but no evidence supports H itself.

`certainly-not:` the negation of H is certain, so H is impossible.

When a proposition 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` or upgrade it to a level higher than `possible.`

Rules are mostly supplied by the user, though there are some built into ATT-Meta. The set of rules is constant through a given application of ATT-Meta. User-supplied rules couch the domain knowledge, knowledge of metaphor and the logical forms for the discourse chunk as mentioned above. In the illustrative syntax used in this report, a simple rule about birds could be

```
IF is-bird(X) AND NOT(dead(X)) THEN {presumed} can-fly(X).
```

A rule's IF-part is usually a conjunction of atomic conditions, or negations of atomic conditions, where atomic means that there are no internal connectives like AND and OR, or use of logical quantification, and the THEN-part is usually an atomic condition or a negation of one. Each rule has a certainty qualifier, which is always one of `suggested, presumed` or `certain.` The level of certainty contributed by the rule to its conclusion when it fires is the minimum of the rule's own level together with the levels of the propositions used by the conjuncts in the IF part. Thus, the rule's qualifier acts as a upper limit on the contribution of the rule.

Rules with a special null IF part count as facts. We will write a fact without any IF or THEN part, as in

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

Facts can have any of the available certainty levels, but in our current practice they are generally `certain.`

Because the particular example dealt with below concerns the possible beliefs of an agent, we now outline how beliefs are represented in ATT-Meta. In the illustrative syntax in this report, a proposition about a belief has the form

`believes(`α`, `λ`, `ϕ`)`

where α is term standing for an agent, λ is a certainty qualifier other than `certainly-not`, and ϕ is the believed proposition. The qualifier λ, which we often call the inner qualifier, is a *lower bound* on the agent's own alleged level of certainty. Thus, the above form can be read as

> α believes, to at least certainty level λ, that ϕ holds.

Of course, a belief proposition `believes(`α`, `λ`, `ϕ`)` also has its own level of certainty, which is independent of the inner level.

The value `suggested` for λ will be particularly important in this report. With this value of λ, the above form can be read as

> α has some reason to believe ϕ.

This does not mean we are entitled to say that α believes ϕ. That statement should be represented by means of an inner qualifier of at least `presumed`.

4 The Original Example and Our Overall Approach to It

4.1 The Example

The utterance of interest in this section is the sentence "*Of course you do,* one small voice insisted" in the following excerpt, (1). We applied the system only to a simplified variant of this sentence, as discussed in the next section. However, it is useful to show in the present section how the original sentence could be handled by our general approach, partly in the spirit of revealing respects in which the ATT-Meta system falls short of it.

(1) Suddenly I was having second thoughts. About us, I mean. Did I really want to get married and spend the rest of my life with Mick? *Of course you do,* one small voice insisted. *Are you quite sure about that?* another nudged. So much was going on in my head, I couldn't sleep.[2]

Notice that the "you" in that example arguably refers to the conscious self of the whole agent.

We assume that the understander, on meeting the "Of course" sentence in (1), has already presumed that another metaphorical view called IDEAS AS INTERNAL UTTERANCES is in play. So the understander presumes that the small voice is an *internal* voice, not a real one. This proposition is placed, according to our approach, inside the pretence cocoon. Then, using knowledge about the source domain of verbal interaction, the understander can infer within the pretence cocoon that there is a person S whose voice it is, and since the voice is inside the writer (the "I" in the discourse) then presumably S is too. Accordingly, as part of within-pretence-cocoon reasoning itself, the understander can work out that MIND PARTS AS PERSONS is in play.

Thus this example is a less obvious case of MIND PARTS AS PERSONS than examples like "One part of me wants to skip the meeting," which are common in mundane discourse.

4.2 Map-Extension Minimization, Mapping Relationships and Ancillary Assumptions

Our application of Map-Extension Minimization to uses of MIND PARTS AS PERSONS lies mainly in our contention *there is no need to try to find some real aspects of the agent's mind that correspond to mentioned or implied "parts,"* and there is no need even to view the mind objectively as being a system with abstract components that entertain different thoughts. That is, although there is *some* analogy operating between the source domain of groups of people and the target domain of the mind, it need not be as strong as to require a componential view of the mind at all, let alone one where individual components can have mental states, and let alone one supporting the identification or postulation of a particular component in a particular case. (1) further illustrates the Map-Extension Minimization default, as we claim that there is no need either to map the notion of "insisting" or of a "small voice" to the mental domain or to map the particular small voice mentioned.

Rather, we only need to assume that the understander, U, has the following two mapping relationships (2,3) and two "ancillary assumptions" (4,5):

(2) The proposition that a subperson S of an agent A *believes,* desires, intends, ... a proposition P (to a high level) maps to the proposition that A *has some reason to believe,* desire, intend, ... P.

(3) An inner conversation among subpersons of an agent A maps, as a whole entity in its own right, to some aspect of A's thinking.[3]

(4) Any utterances by a subperson S of an agent A are part of a *private conversation* involving only his/her subpersons (one of which can be the conscious self of A).

This means in particular that

- if a subperson S of an agent A utters something to someone O, then O is a subperson of A

[2] *My Story* magazine, May 1995, Editions Press Limited, Gibraltar, p.6/7. Italics in original. Paragraph breaks deleted. The example is contained in the MIND PARTS AS PERSONS (OR OTHER ANIMATE BEINGS) section of our databank, Barnden (n.d.).

[3] We do not specify what this aspect is in this report, and the mapping need not in principle do so—it could just postulate the existence of *some* corresponding aspect of A's thinking.

– any utterance by a subperson S of an agent A is (by default) heard by all other subpersons of A, and is not heard by any other agent.

(5) The conversation in (4) is *physically inside* A.

Ancillary assumptions such as (4,5) are proposed in our approach as an important aspect of some metaphorical views, over and above mappings between the source and target domains. Ancillary assumptions serve to enrich the nature of the assumed source-domain scenarios in uses of the metaphorical view.

4.3 Connection to Some Other Metaphorical Views

Because of the second ancillary assumption (5) and the second mapping relationship (3), MIND PARTS AS PERSONS examples that involve utterances by the mind-parts are also manifestations of IDEAS AS INTERNAL UTTERANCES, an important metaphorical view that we will not address extensively in this article but that is discussed in Barnden (1997).

It is plausible also that MIND PARTS AS PERSONS is a special case of MIND AS PHYSICAL SPACE, with the mind-parts being seen as being in the mind-space. Indeed, this is suggested by the clause "So much was going on in my head" in the example, if the mind-space is viewed as overlapping the interior of the head. This additional possibility will not be pursued further in this paper, as the example can be handled without considering it.

MIND PARTS AS PERSONS overlaps IDEAS AS ANIMATE BEINGS in applications of the latter that involve the ideas being viewed as physically located within the agent. However, IDEAS AS ANIMATE BEINGS usually focuses on non-verbal physical behaviour of the ideas/beings rather on their own mental states or intercommunications. The idea can be portrayed, for instance, as "lurking."

4.4 Main Reasoning Involved in the Example

The overall rough pattern of reasoning that we suggest for the example is sketched in Figure 2. We noted above that the understander infers within the cocoon that there is a person S physically inside Mary that owns the mentioned small voice (SV in the Figure). Similarly, the owner of the second mentioned voice (OV in the Figure) is presumably also a subperson O of Mary. Now, the fact that S *insists* something implies that S asserts it, which in turn implies (by default) that S believes it. The fact that the second subperson, O, verbally questions S's assertion implies (by default) that O has a *mental* questioning attitude towards what was asserted by S.

These inferences are entirely performed using *ordinary commonsense knowledge about speech and thought.* There is nothing specific to metaphor or to any metaphorical view in this reasoning. The first ancillary assumption (4) can be used, however, to provide additional support for the proposition that subperson S is talking to O, the owner of the other voice. In fact, O is presumably the conscious self of the "I," because of the first three sentences of (1).

Let us use A to denote the whole, real agent—the persona uttering (1) and referred to by "I" there. As a result partly of inferring within the cocoon that S believes that A wants to marry Mick and that O mentally questions whether A wants to marry Mick, the understander can apply the mapping relationship (2) to both S and O, separately. By this means the understander can infer (by default) the target-domain informational contributions that

A has some reason to believe that she wants to marry Mick

and that

she has some reason to question whether she wants to marry Mick.

4.5 More Subtle Connotations

We now attend to the significance of the adjective "small" qualifying voice SV in (1), and to an additional implication of the voice *insisting* something as opposed to merely asserting something. See Figure 3. Suppose that the understander can infer from the smallness that the voice's owner S is relatively unimportant in the inner conversation inside A. The understander can then infer that the fact that S believes that A wants to marry Mick is relatively unimportant in the inner conversation. This is a value-judgement by the understander on a within-pretence state of affairs *that maps to a state of affairs in reality.* The latter, real state of affairs is that of A having some reason to believe that she wants to marry Mick. Hence, by the Value-Judgment VNMA, the understander can infer that this state of affairs is relatively unimportant in A's thinking. Since no such consequence is derived for A's having some reason to mentally question whether she wants to marry Mick, it follows that *A's having some reason to mentally question the marriage desire is more important than A's having some reason to believe it.*

However, the insistence by S implies (defeasibly) that the proposition that the *issue of whether* A wants to marry Mick is of great importance *to S*. It is plausible to suggest that the understander could infer from this that that issue is of great importance within the inner conversation among the subpersons, not just to S. The understander could then use the Value-Judgment VNMA to infer that *the issue is also important within A's thinking,* given that the inner conversation maps to an aspect of that thinking by (3).

Thus the approach is able to deal with both of two subtly different value judgments conveyed by the metaphorical utterance:

- A's *having reason to believe* that A wants to marry Mick is less important in A's thinking than A's having reason to question that desire;
- the *issue of whether* A wants to marry Mick is, nevertheless, important to A.

4.6 Stockness and Sidelining

We stated above that we are mainly concerned with metaphorical utterances that do not just use stock phraseology and whose metaphoricity is not sidelined. Now, speech acts are commonly attributed to subpersons within people, and voices within people are commonly mentioned. The online Cambridge International Dictionary of English (CIDE)[4] lists the following examples of an inner voice uttering something:

> A voice within you is your conscience the part of you which tells you when you are doing something immoral. [lack of punctuation is as in original]
>
> A voice within him said he shouldn't be driving so fast.
>
> Suddenly this voice within her told her to stop being so stupid.

[4] `http://dictionary.cambridge.org`

Moving closer to (1), a search over the full Bank of English (a corpus of 415 million words)[5] gave us (only) 50 examples of "small voice" in a context implying that the voice is within someone.[6]

Of the 50 instances of inner "small voice," 23 were accompanied by some other speech verb attributing speech to the voice. The maximum number of word in between was seven. However, we only found 4 examples in the Bank of English of the word "voice" accompanied by any form of the verb "insist" within a separation of seven words and where the insisting was attributed to the voice and the voice was inner. None of those examples classified the voice as "small." We therefore feel justified in not regarding our example as being just a matter of stock phraseology.

As for sidelining, it is hard to think of senses of the adjective "small", noun "voice" and the verb "insist" that would conspire to provide the set of contributions to the overall meaning of the example sentence that the analysis above provides. Certainly, the entries in two major dictionaries—CIDE and Webster's Third New International Dictionary—provide no such senses. The closest approach in those entries is that Webster's does give the sense "persist in a point of view about" for the verb "insist," and this sense description could be interpreted as having a literal meaning in the mental domain (although the example given in Webster's is in fact about speech).

We therefore feel justified in taking the metaphoricity of the sentence not to be sidelined. Of course, a given human or artificial understander could differ from published dictionaries on what lexical senses words have, and this point is an example of our understander-relativity of metaphor and in particular of sidelining (see the Introduction above).

Also, a lexicon could list some such template as "[ADJECTIVE] voice [SPEECH VERB]" and assign it a sense directly in the mental domain. However, this sense would need to take account of all the possibilities for the adjective and all the different connotations of different speech verbs. The sense would in addition need to bring in the implicit conflict that is inherent in MIND PARTS AS PERSONS examples, because of there often being more than one subperson in the implied metaphorical scenario.

5 Altered Example and Its Treatment in the ATT-Meta System

Our application of the ATT-Meta system was to the following variant of part of (1).

(6) A voice inside Mary insisted that Mick was adorable.

We altered the example because:

- it is simpler to discuss the example if the voice makes a simple claim about Mike rather than a claim about Mary wanting to marry Mike;
- we have not handled the smallness of the voice in our implementation, and indeed the system does not yet have a systematic treatment of the Value-Judgment VNMA needed to handle the smallness;
- and we have no proper handling of mental questioning in the system.

[5] `http://www.cobuild.collins.co.uk/boe_info.html`

[6] Of these, 26 implied that the voice was the voice of God or some (other) inner spiritual aspect of the person; almost but two of these examples were in the idiomatic form "still[,] small voice." That phrase was also used 6 times to refer to an ordinary, non-spiritual mental event. The phrase "small voices" only appeared 6 times in the Bank of English, so we did not consider the plural form further.

We further simplify as follows. We assume that when one person insists in a conversation that something is the case, presumably someone else in the conversation verbally *denies* that it is the case. We therefore ignore the possibility that the someone-else merely questioned it or ignored it, which would be other reasons for the first person to insist. In addition, our implementation did not deal with the importance that insistence confers upon what is insisted as far as the insister is concerned.

We also assume, partly for simplicity, that in the mental "conversation" in Mary's mind there is just one other person other than the owner of the mentioned voice. This is also partly to get round the fact that ATT-Meta currently does not fully support quantification within queries and propositions. Ideally, we want to say that *some* internal interlocutor denies what the voice says.

As against these simplifications, there is a sense in which we made the task more difficult for ATT-Meta. In discussing (1) we assumed that the understander extracts from the discourse the fact that there is another inner subperson (owner of the voice that "nudged"). However, in dealing with (6) we left it up to the system itself to infer, from the presence of one subperson, the presence of another subperson, who is the interlocutor of the owner of the explicitly mentioned voice.

Figure 4 shows some of reasoning involved in inferring the informational contribution that Mary had reason to believe that Mick was adorable and also had reason to believe that he wasn't. The following subsections go into the detail of how the reasoning arises.

5.1 User-Supplied Information about Specific Situation

The information we supplied about the specific situation to be reasoned about was as follows. We imagine the information to arise by preliminary processing of the example sentence together with the results of understanding prior context.

```
{certain} is-person(Mary)

IF  currently-within-metaphor-pretence
AND IN SURROUNDING CONTEXT: true
THEN {certain} is-voice(v).

IF  currently-within-metaphor-pretence
AND IN SURROUNDING CONTEXT: true
THEN {certain} says(v, adorable(Mick), insistingly).

IF  currently-within-metaphor-pretence
AND IN SURROUNDING CONTEXT: true
THEN {certain} physically-in(says(v, adorable(Mick), insistingly),
                             mary).
```

Here `v` is a constant symbol standing for the voice. In accordance with one of our simplifications above, we assume that the owner of the voice has a unique interlocutor, given by the function `interlocutor-of`. Notice, however, that the interlocutor itself is *not* given as being physically inside Mary. That proposition will arise by inference.

The above rules, aside from the first, are designed for application only to queries within the pretence cocoon, as shown by the special condition `currently-within-metaphor-pretence`. In general,

the conditions in such a rule establish subqueries within the pretence cocoon, except for conditions prefixed by "IN SURROUNDING CONTEXT: " which establish subqueries in the next reasoning context out, in other words the reality space. In the present case those subqueries will be just the vacuous query `true`, which is always supported to full certainty. Overall, the effect of the two rules is to ensure that, within the pretence cocoon, the queries

```
is-voice(v)

says(v, adorable(mick), insistingly)

physically-in(says(v, adorable(Mick), insistingly), mary)

is_interlocutor-of(interlocutor-of(owner_of(v)), owner_of(v))
```

will be supported to level `certain`.

5.2 User-Supplied Query

We supplied the following top query to the system:

```
(7)
  believes(Mary, suggested, adorable(Mick))
  AND
  believes(mary, suggested, NOT(adorable(Mick))).
```

Intuitively, bearing in mind the comments on belief representation at the end of section 3, the query was: Does Mary have reason to believe Mick is adorable and reason to believe that he isn't? The system gives an affirmative answer, assigning `presumed` as the final certainty level to the above query.

However, we do not assert that a query as precise as the above would normally arise from context. Rather, the query might be on the lines of "What is Mary's mental state concerning Mick?" Also, the content of the sentence itself could suggest that the relevant issue is whether Mick is adorable. The sophisticated processing of context implied here is not something we have implemented in ATT-Meta. A more technical limitation is that ATT-Meta cannot currently handle a query like "What does Mary believe about Mick?"

5.3 User-Supplied Rules about Mental States and Physical Objects/Space

Various rules about mental states and about physical objects/space in general were supplied to the system. They were the same ones as were supplied for the Anne/Kyle application in (Barnden & Lee 2001) and for the applications in Lee & Barnden (2001b). However, in the present application none are *strongly-used*—none delivered a certainty contribution of at least `presumed`)—so we do not show them here.

5.4 User-Supplied Rules about Verbal Communication

The following are the strongly-used rules about verbal communication that were supplied to the system.

```
IF is-voice(V)
THEN {presumed} is-person(owner-of(V)).

IF is-person(P) AND is-interlocutor-of(I,P)
THEN {presumed} is-person(I).

IF is-voice(V) AND says(V, Proposition, Manner)
THEN {certain} says(owner-of(voice), Proposition, Manner).

IF is-voice(V) AND physically-in(V,R)
THEN {presumed} physically-in(owner-of(V),R).

IF is-voice(V) AND says(V, Proposition, Manner)
AND physically-in(says(V, Proposition, Manner), R)
THEN {presumed} physically-in(V,R).

IF says(P, X, insistingly) AND is-person(P)
AND is-interlocutor-of(I,P)
THEN {presumed} says(I, NOT(X), ordinary_manner)

IF says(P, Proposition, Manner) AND is-positive-manner(Manner)
AND is-person(P)
THEN {presumed} believes(P, presumed, Proposition)

{certain} is-positive-manner(ordinary-manner).

{certain} is-positive-manner(insistingly).
```

5.5 Rules for Metaphorical Views

We provided to the system the same rules for metaphorical views as we did in the Anne/Kyle example in (Barnden & Lee 2001). These included rules for MIND PARTS AS PERSONS, of course, but also rules for IDEAS AS PHYSICAL OBJECTS and MIND AS PHYSICAL SPACE. However, the rules for the latter two views were not strongly used in the reasoning for the current example, so we present here only the strongly-used rules for MIND PARTS AS PERSONS.

The only mapping relationship for MIND PARTS AS PERSONS that we needed for dealing with (6) is one from an inner person having a belief, with certainty at least `suggested` within that inner person, to the whole agent having the belief, with certainty at least `suggested` within the agent. Cf. mapping relationship (2) above. We did not include an implementation of mapping relationship (3) above, which is not needed for the reasoning addressed in the present section.

If the rule corresponding to (2) fires it merely establishes that the whole agent believes the belief to certainty (at least) `suggested`, no matter what the inner certainty of the inner person is. So, in our example, even if the inner person were `certain` that Mary wants to marry Mick, Mary would still only be

inferred to have a certainty of (at least) `suggested` that she wants to marry Mick. Of course, this leaves it *open* that she has a higher certainty.

We also included an ancillary-assumption rule for MIND PARTS AS PERSONS saying that any interlocutor of an inner person is, by default, (physically) in the agent as well. This rule corresponds to local assumption (5) above. We also partially implemented (4) by including a fact rule that has the effect of postulating an unique interlocutor of the owner of the voice mentioned in (6). We did not implement the other aspects of (4), which are not needed for the reasoning in this section.

The conversion rule and ancillary assumption rules just mentioned are respectively as follows.

```
(8)
   IF currently-within-reality-context
   AND is-person(A)
   AND within pretence: is-person(X)
   AND within pretence: physically-in(X,A)
   AND within-pretence: believes(X, suggested, Proposition)

   THEN {presumed} believes(A, suggested, Proposition)

(9)
   IF currently-within-metaphor-pretence
   AND is-person(A)
   AND is-interlocutor-of(X,Y)
   AND is-person(Y)
   AND physically-in(Y,A)

   THEN {presumed} physically-in(X,A).

(10)
   IF  currently-within-metaphor-pretence
   AND IN SURROUNDING CONTEXT: true
   THEN {certain} is_interlocutor-of(interlocutor-of(owner_of(v)),
                                     owner_of(v)).
```

Rule (9) would not be needed if there were a general default rule about real verbal communication that said that any interlocutor of someone who is physically in a region X is also in region X.

5.6 Sketch of Sequence of Events

The conjunctive query (7) in section 5.2 leads to the conjuncts being posed as separate subqueries. The THEN part of (8) matches the first subquery. Therefore the conditions in (8) generate the within-pretence subqueries

```
   is-person(X)
   physically-in(X, Mary)
   believes(X, suggested, adorable(Mick)).
```

These are straightforwardly satisfied inside the cocoon, with X bound to `owner-of(v)`, by virtue of the above pretence facts and common-sense rules and facts about people and verbal communication. See Figure 4.

The other conjunct of the top-level query (7) also matches (8) and therefore leads to the within-cocoon subqueries

```
is-person(X)
physically-in(X, Mary)
NOT(believes(X, presumed, adorable(Mick))).
```

These also become satisfied, with the subperson variable `X` now bound to `interlocutor-of(owner-of(v))`. This is largely by virtue of the pretence facts and common-sense rules and facts about people and verbal communication. But the reasoning also uses ancillary assumption rules (9) and (10) and a rule that infers

`NOT(believes(`α`, presumed,` ϕ`))`

from

`believes(`α`, presumed, NOT(`ϕ`))`

for any α and ϕ. This last rule is a "belief management" rule built into ATT-Meta. The ancillary-assumption rule (9) allows ATT-Meta to infer (by the link marked AA in Figure 4) that `interlocutor-of(owner-of(v))` is physically in Mary, a condition required for (8) to be able to fire in the case of the interlocutor.

5.7 Some Computational Statistics

The set of user rules in the implementation included metaphorical mapping and ancillary-assumption rules for IDEAS AS PHYSICAL OBJECTS and MIND AS PHYSICAL SPACE as well as MIND PARTS AS PERSONS, and mini knowledge bases for the domains of verbal interaction, physical space/objects and mental states. As a result, there were 63 user rules not counting the specific-situation facts, of which there were 5. In addition, at the time of experimentation ATT-Meta had 27 built-in rules, dealing with qualitative degrees, metaphorical pretence in general, belief states in general, and miscellaneous general inferential matters.

In the example of section 5, the system created 162 hypotheses in the course of reasoning. Of these, 31 attained a certainty level of `presumed` and 17 attained *certain.*

Of the total 162, two were in the belief space of Mary and 141 were within the pretence cocoon (either directly or within the belief spaces of the two subpersons). 25 of the 141 ended up at certainty level `presumed` and 11 at *certain.*

Without any user rules at all, 4 hypotheses are generated, so that the 17 strongly-used user rules (including specific-situation facts) added 70 hypotheses, giving an average of 4.12 hypotheses per strongly-used rule.

5.8 Rule-Base Overhead and Scale-Up

Partly in order to check the overhead resulting from the inclusion of extraneous rules, we also ran the system after removing all except the 17 *strongly-used* user rules (including facts): recall that a rule is strongly used if at some point it delivers a certainty contribution of at least `presumed` to some hypothesis. (All built-in

rules were kept even if they were not strongly used.) As a result, the number of hypotheses created dropped to 74. Of these, 38 attained a certainty level at least `presumed.` Of the total of 74, 51 were within the pretence cocoon, 26 of them at level at least `presumed`.

Comparing the full run with the slimmed-down run, the 51 non-strongly-used user rules added 88 hypotheses, i.e. an average of about 1.73 hypotheses per non-strongly-used rule. Thus, the more irrelevant rules add only a relatively small number of hypotheses per rule (compare with the 4.12 hypotheses per strongly-used rule, noted above).

As a further test of the effects of including more rules, we added the additional rules used in the various implemented examples in Lee & Barnden (2001b). The domain rules amongst these rules concerned possession of objects, fashion, resource usage, life and death, vision, physical constructions and personal relationships. The additional metaphorical-view rules (i.e., conversion rules and ancillary assumptions) concerned views of beliefs as possessions, mental operation as resource usage, idea as animate beings, cognizing as seeing, ideas as physical constructions and beliefs as locations. The effect of the 127 added rules was to increase the number of hypotheses generated from 162 to 285. The increase was therefore at a rate of only 0.97 of a hypothesis per rule.

These observations are encouraging, in that many of the additional rules were on subject matter strongly related to that of the original rule-set, and the added metaphorical mapping rules were related to the original ones in also being about mental states. However, we do not yet known how the effect will scale up with larger numbers of domains, rules per domain and metaphorical views.

One reason for caution is that there can be interactions between different sets of rules. For example, when only the eight MIND AS PHYSICAL SPACE rules are removed from the full original set of rules (which led to 162 hypotheses), the hypotheses dropped to 125 (a reduction of 4.63 hypotheses per rule excluded). When those eight rules are removed from the implementation augmented with the additional rules from Lee & Barnden (2001b), the hypotheses dropped more sharply from 285 to 217 (a reduction of 8.50 hypotheses per rule excluded).

6 Conclusion

The report has shown that the ATT-Meta system can derive what are arguably the central informational contributions from a variant (6) of a real-discourse metaphor example (1) based on a commonly-used metaphorical view, MIND PARTS AS PERSONS. The contributions are the two conjuncts of the user query (7) in section 5.2, once these contributions have become `presumed`. These contributions are that Mary has reason to believe that Mick is adorable and reason to believe that he is not. The example was a mundane, non-sidelined, map-transcending utterance. Such utterances are the main target of the ATT-Meta approach.

In particular, the report gives suggestive grounds for supposing that the ATT-Meta approach's Map-Extension Minimization is on the right track. The information extracted from the example utterance was extracted in the absence of any map-extension, even though the mapping that we assumed to be involved in MIND PARTS AS PERSONS is extremely limited (see (2) in section 4 and (8) in section 5.5). In particular, the existing mapping did not need to be extended to map the voices, subpersons and subperson speech act-types (e.g., insistence) implied by the example.

The current implementation falls short of what we would ideally like, as it does not deal with the implications of the smallness mentioned in the original example (1), nor does it deal with the importance conferred by the verb "insist" on what is insisted. In order for such relatively subtle matters to be dealt with in an non-ad-hoc manner, we would need an implementation of the Value-Judgment VNMA. However, VNMAs are not fully implemented in ATT-Meta. Some VNMAs are currently implemented or implementable in ATT-Meta in limited or ad hoc ways. Notably, for any given conversion rule that we implement, we

normally accompany it by a negative version, thus getting the effect of the Negation VNMA. (The MIND PARTS AS PERSONS conversion rule (8) is exceptional in that it does not have a negative version in our implementations.) Similarly, many conversion rules that we implement handle qualitative degrees explicitly, so that the effect of the Qualitative Degree VNMA is achieved. However, these implementations leave a lot to be desired as they require the user to explicitly supply extra rules or rule-parts instead of allowing them to be automatically supplied by the system.

Acknowledgments

The research is being supported by grant GR/M64208 from the Engineering and Physical Sciences Research Council (UK).

References

Barnden, J.A. (1997). Consciousness and common-sense metaphors of mind. In S. O'Nuallain, P. McKevitt & E. Mac Aogain (Eds), *Two Sciences of Mind: Readings in Cognitive Science And Consciousness*, pp.311–340. Amsterdam/Philadelphia: John Benjamins.

Barnden, J.A. (1998a). An AI system for metaphorical reasoning about mental states in discourse. In Koenig, J-P. (Ed.), *Discourse and Cognition: Bridging the Gap*, pp.167–188. Stanford, CA: CSLI.

Barnden, J.A. (1998b). 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. (2001). Application of the ATT-Meta metaphor-understanding approach to selected examples from Goatly. Technical Report CSRP–01–01, School of Computer Science, The University of Birmingham, U.K.

Barnden, J.A. (n.d.). ATT-Meta Project Databank. `http://www.cs.bham.ac.uk/~jab/ATT-Meta/Databank`.

Barnden, J.A. (in press). The utility of reversed transfers in metaphor. To appear in *Procs. Twenty-Third Annual Meeting of the Cognitive Science Society* (Edinburgh, August 2001). 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.

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), *Second International and Interdisciplinary Conference on Modeling and Using Context* (CONTEXT'99), *Lecture Notes in Artificial Intelligence, 1688*, pp.28–41, Springer.

Barnden, J.A. & Lee, M.G. (2001). Understanding open-ended usages of familiar conceptual metaphors: An approach and artificial intelligence system. Technical Report CSRP–01–05, School of Computer Science, The University of Birmingham, U.K.

Goatly, A. (1997). *The language of metaphors.* London and New York: Routledge.

Lakoff, G. (1993). The contemporary theory of metaphor. In A. Ortony (Ed.), *Metaphor and Thought*, 2nd edition. Cambridge University Press.

Lakoff, G. (1994). The Master Metaphor List. `http://cogsci.berkeley.edu`. University of California, Berkeley.

Lakoff, G. (1996). Sorry, I'm Not Myself Today: The metaphor system for conceptualizing the self. In G. Fauconnier & E. Sweetser (Eds), *Spaces, Worlds, and Grammar*, pp.91–123. Chicago, IL: University of Chicago Press.

Lee, M.G. & Barnden, J.A. (2001a). Reasoning about mixed metaphors with an implemented AI system. *Metaphor and Symbol, 16* (1&2), pp.29–42. Also as Technical Report CSRP-00-11, School of Computer Science, The University of Birmingham, U.K.

Lee, M.G. & Barnden, J.A. (2001b). Mental metaphors from the Master Metaphor List: Empirical examples and the application of the ATT-Meta system. Technical Report CSRP–01–03, School of Computer Science, The University of Birmingham, U.K.

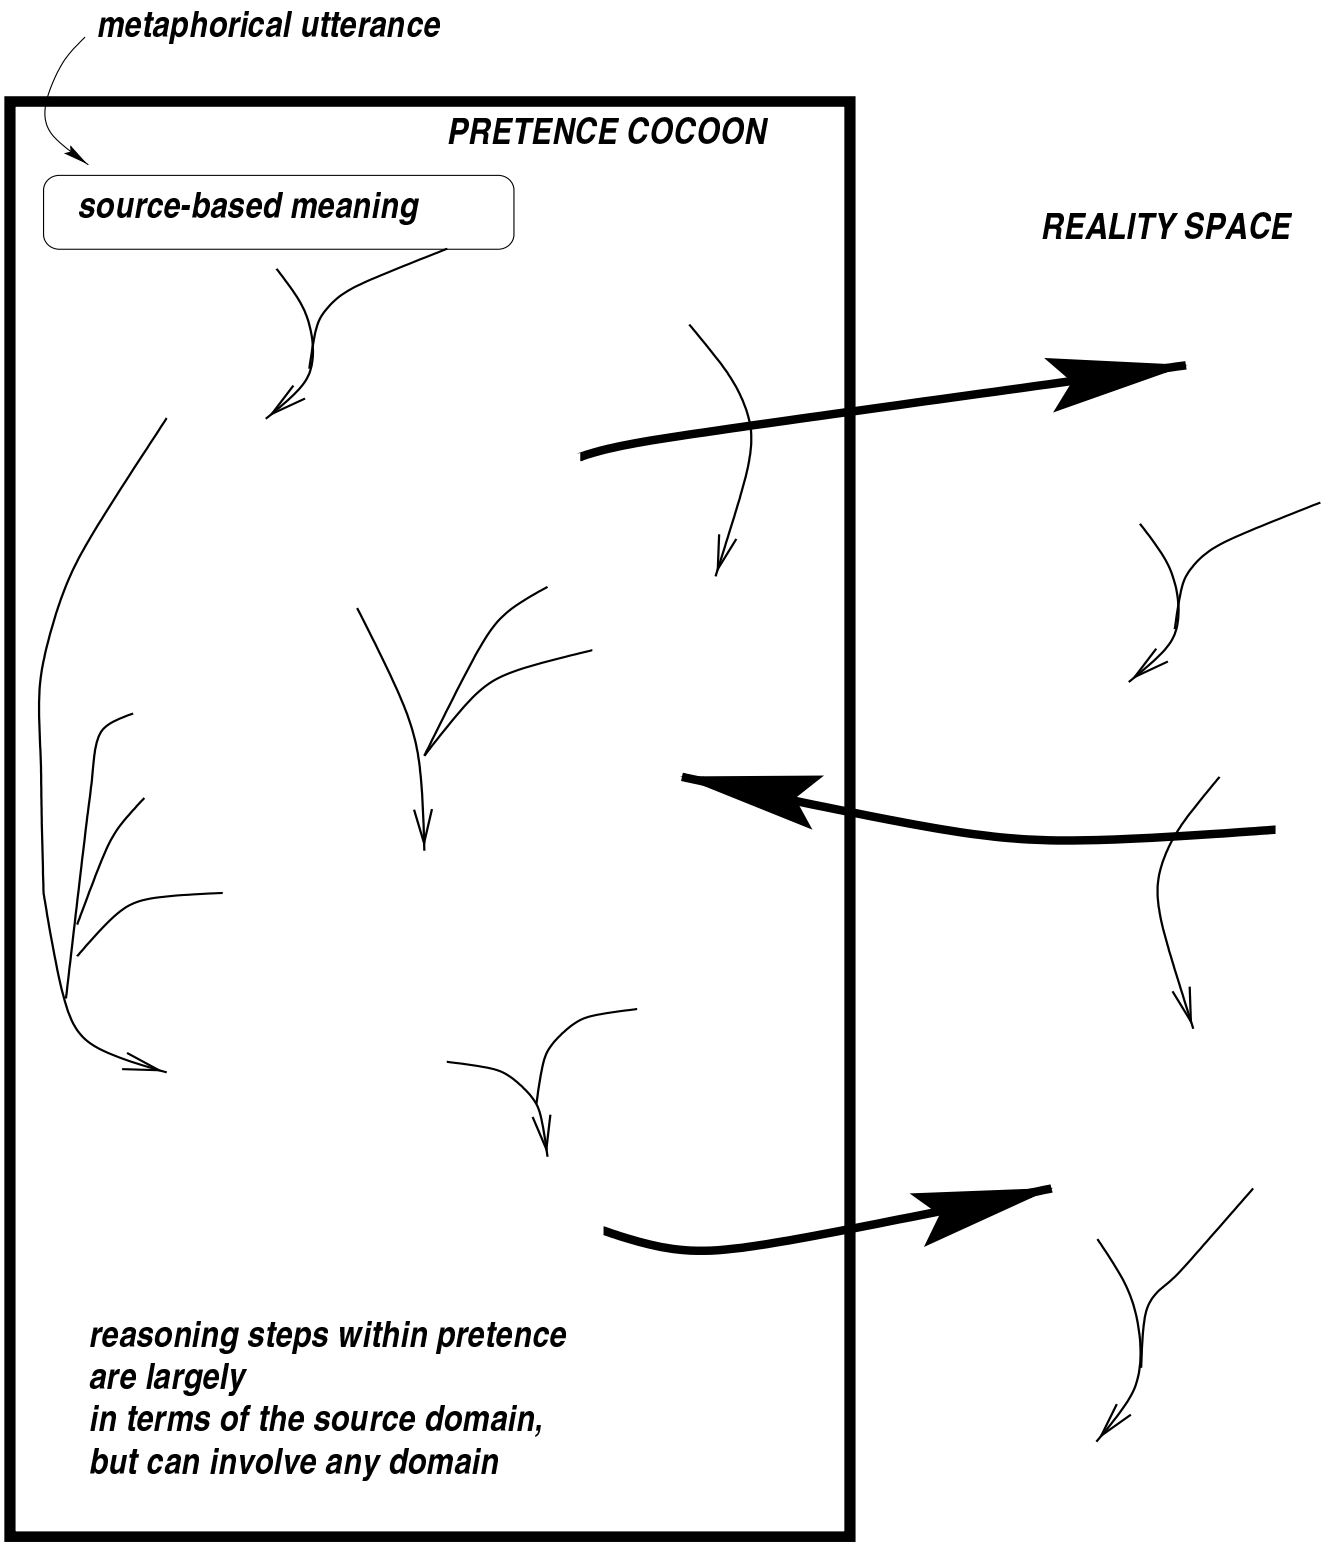

Figure 1: **General nature of reasoning in the ATT-Meta approach.** The bold box shows the pretence cocoon. Bold arrows show the action of mapping relationships, between source-domain information in the pretence cocoon and target-domain information in reality. (Target-to-source mapping actions are allowed in the approach, for reasons explained in Barnden, in press. An example occurs in Barnden (2001: Figure 10).) Other arrows, apart from the one at the top left, show reasoning actions within reality or within the pretence.

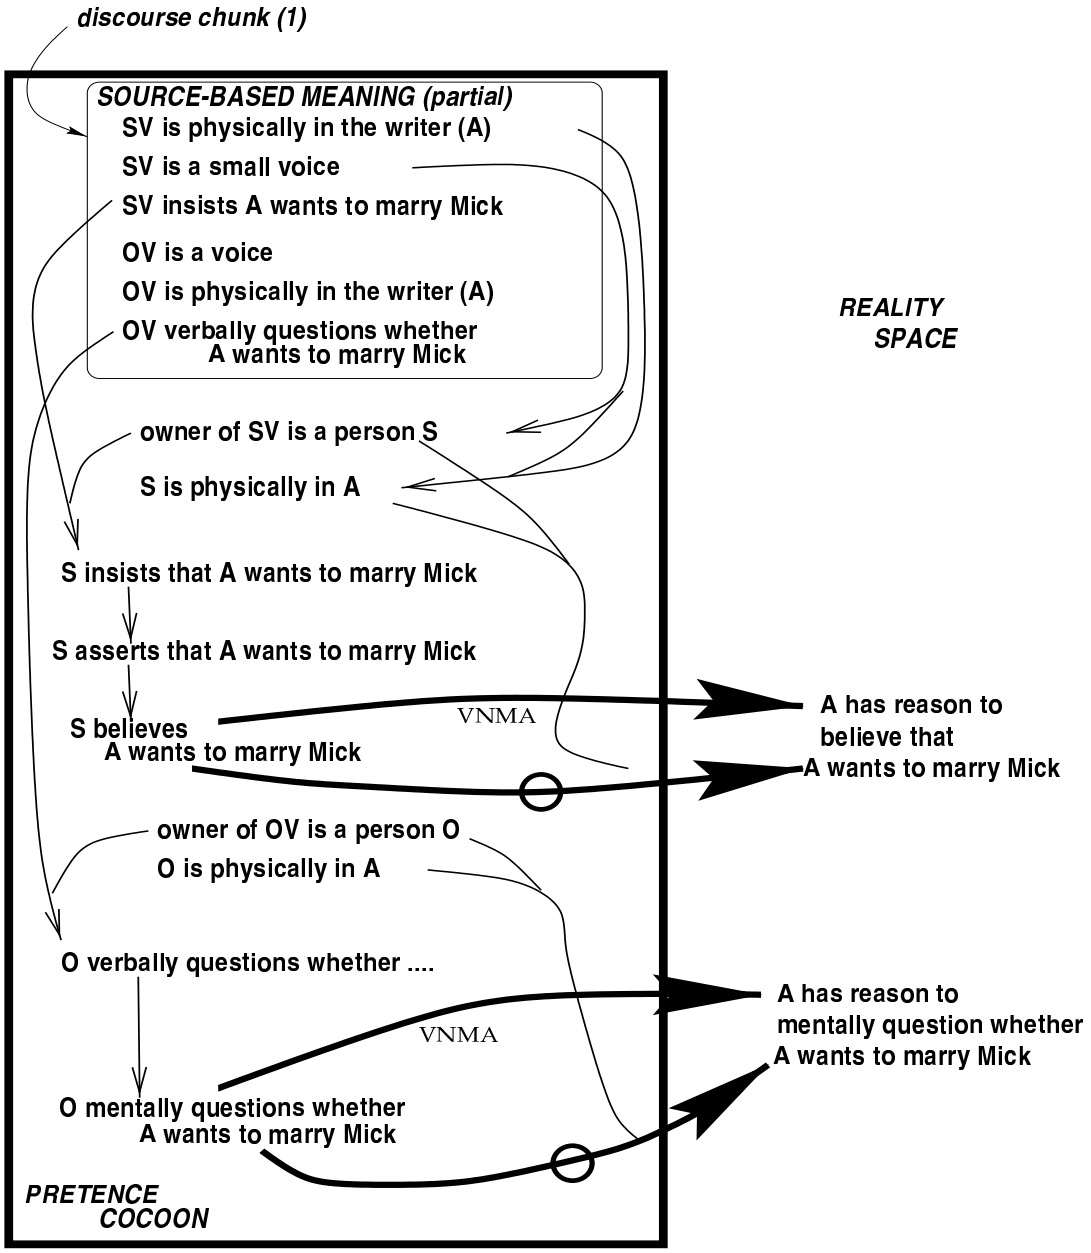

Figure 2: In this and later figures, a mapping arrow labelled VNMA or VNMAs shows the action of one or more view-neutral mapping adjuncts. A mapping arrow marked with a circle shows the action of a mapping relationship specific to a particular metaphorical view, here MIND PARTS AS PERSONS. The source-based meaning of the utterance is shown in a small box near the top of the diagram. The propositions within the diagram are English glosses of expressions in some internal representation scheme used by the understander. In the figures only a selection of the possible propositions and inferential links are shown. Tense in the English glosses is actually handled by the Time-Order VNMA (Barnden & Lee 2001), together with identity mapping of the present time point. Therefore, the mapping arrows marked by a circle as view-specific often represent in actuality the combined operation of a view-specific mapping relationship and the Time-Order VNMA.

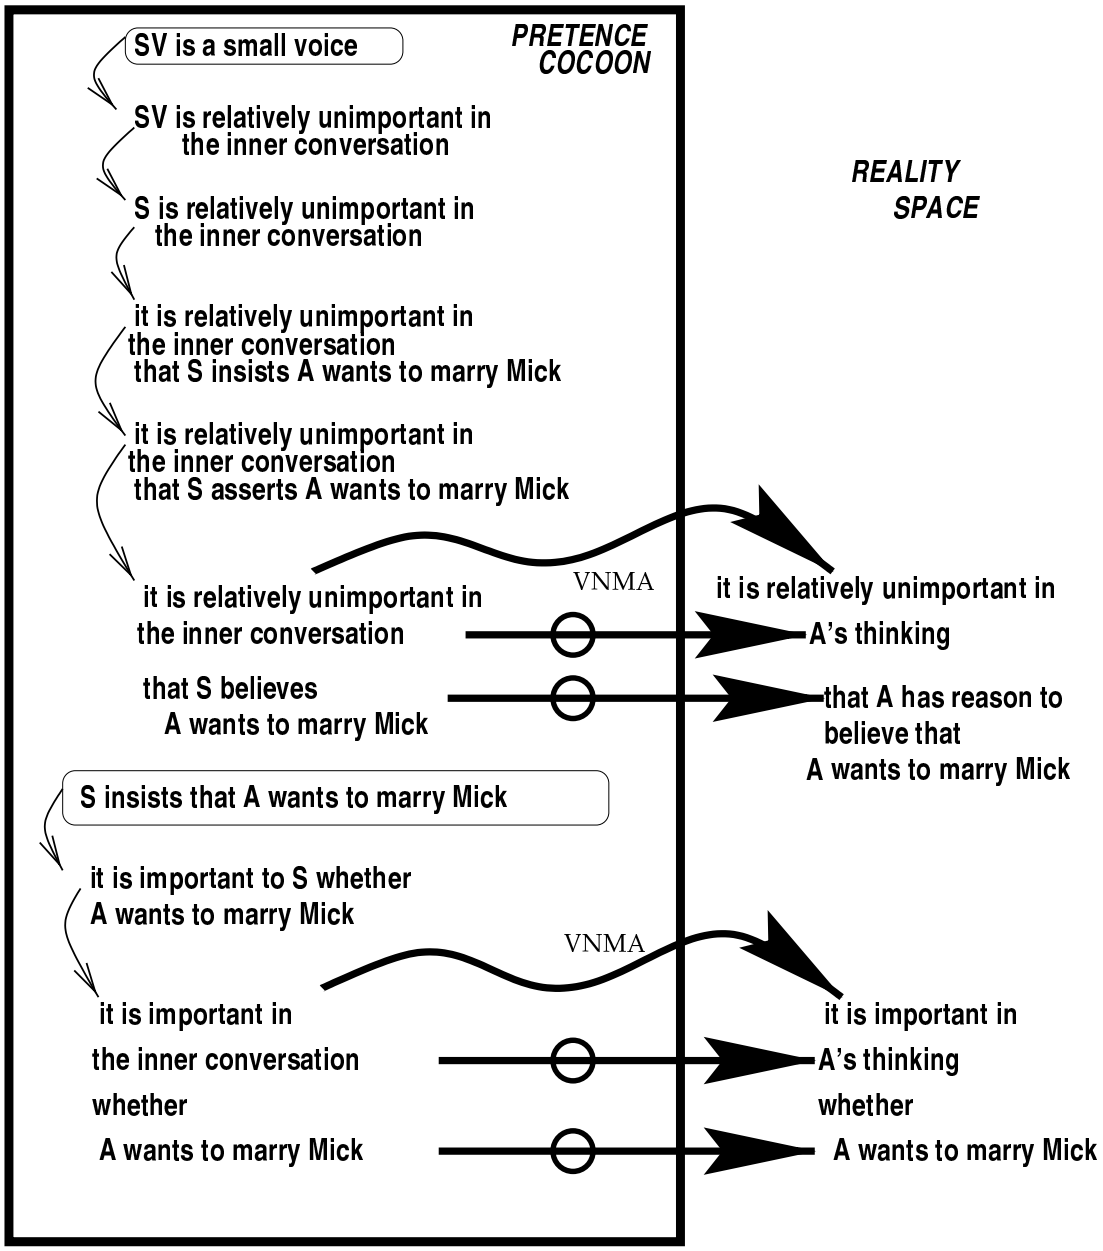

Figure 3: Dealing with the smallness of the voice and the nature of insistence. The VNMA used is the Value-Judgment VNMA. The inner conversation is shown for brevity as mapping to A's thinking, but in fact maps to some aspect of A's thinking.

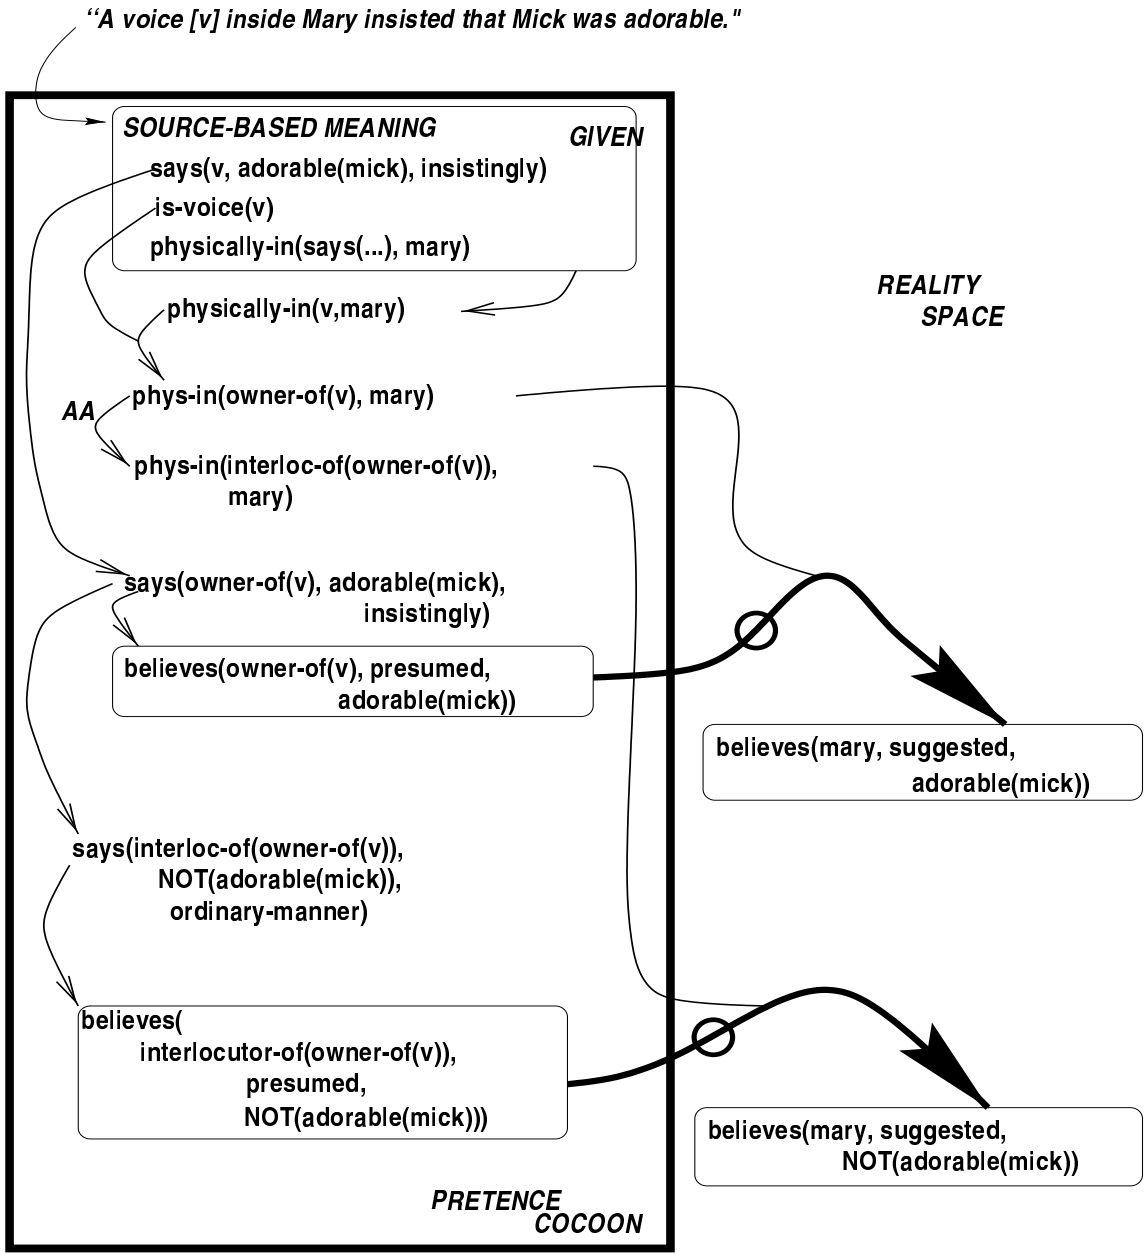

Figure 4: Showing the main features of ATT-Meta's metaphor-based reasoning for the small-voice example (6). Not all hypotheses and rule applications needed to support the depicted hypotheses are shown. The fat arrows depict applications of conversion rule (8). The application of ancillary assumption rule (9) is shown as the link marked AA.