Application of the ATT-Meta metaphor understanding system to examples of the metaphorical view of TEACHERS AS MIDWIVES

(Technical Report CSRP-02-10)

John A. Barnden, Sheila R. Glasbey, Mark G. Lee and Alan M. Wallington
School of Computer Science, The University of Birmingham
Birmingham, B15 2TT, United Kingdom

J.A.Barnden@cs.bham.ac.uk

Abstract

This report is intended as a companion to another report, *Asymmetry and Reverse Transfers in Metaphor*, [Barnden et al., 2002]. In that report, it is argued that the transfer of information and other effects from the tenor or target domain to the vehicle or source domain plays a highly important, albeit subsidiary, role in metaphor understanding; the main transfer in metaphor understanding goes from source to target. Furthermore, it is claimed that the importance of this type of transfer has been underestimated or overlooked entirely, hidden from detailed study by a confusion between different notions of direction and asymmetry in metaphor that one finds in the literature. However, the latter report, whilst it refers to an AI program for metaphorical reasoning called ATT-Meta, which implements target to source (and source to target) transfers, does not report any results from a simulation. It does not give a detailed listing of the rules used in illustrative discourse examples exemplifying the various types of reverse transfer, and it does not give a step-by-step account of ATT-Meta's reasoning as it processes the different examples. This report provides enough detail of the ATT-Meta system to understand its approach to metaphorical reasoning, and it lists the rules and shows the reasoning involved in processing the examples.

0 Introduction

Metaphorical views[1] are understood asymmetrically. A view of teachers as being midwives is not the same as a view of midwives as being teachers. By convention, the first term in the metaphorical view corresponds to the target or tenor and the second term corresponds to the source or vehicle. Therefore, in the case of the metaphorical view of TEACHERS AS MIDWIVES, it is teachers who are viewed as having certain of the characteristics, behaviour and so on of midwives, and not the converse. Somewhat paradoxically the specific mapping links that relate target and source elements are usually described with the source domain first, so going from source to target, or S → T. Thus, the metaphorical view of TEACHERS AS MIDWIVES is usually depicted with a directed arrow going from midwives to teachers, or midwives → teachers.

An easy assumption to make is that the asymmetry of a metaphorical view should correspond to an asymmetry of transfer. In other words, when a metaphor is being read and understood, the processing of the source domain will affect in some manner the understander's view of the target domain by transferring aspects of the source to the target. This is to be contrasted with the situation −which is assumed not to happen− in which the processing of the target domain affects in some manner the understander's view of the source domain, even though the mapping link is S → T. In other words, there is assumed to be directionality in processing. However, it is important to realise that an asymmetry in view does not entail an asymmetry in use. In *Asymmetry and Reverse Transfers in Metaphor*, [Barnden et al., 2002] it is argued that it is important to distinguish between the direction of main transfer in understanding a metaphorical view, and the directions of individual transfers. The former is inherent in the very terms "source" and "target" and we can assume that the main goal of the metaphor is to affect this transfer and so use information from the source to affect the target in some manner. However, whilst still preserving this main goal over a stretch of discourse, it is argued that within a particular metaphorical stretch there may be cases when the processing of the target domain does affect the source domain, which can then in turn affect the target. Furthermore, it is argued that these situations are quite common. In other words, there are situations where "reverse transfers" from the target domain to the source domain.

Arguments supporting the claim that instances of reverse transfer occur in metaphor understanding, together with illustrative examples are given in *Asymmetry and Reverse Transfers in Metaphor*, [Barnden et al., 2002], henceforth ART. However that report, although it briefly discusses how the ideas might be implemented in the ATT-Meta computer system, does not provide a sufficiently detailed account of the rules needed to allow the reader to trace through a particular example. The intent of this report is to provide that detailed account.

We give an account of the application of the ATT-Meta system to the different types of T→ S transfer discussed in ART, listing, using a notational variant of the rules used in the actual ATT-Meta system, the various rules required, and showing how they interact. Our examples are taken from section 4 of that report and involve the running example "SOCRATES AS MIDWIFE" involving 'Socrates', 'Electra', and 'Electra's idea'. This view is extensively developed in Plato's Theaetetus (Plato, 1973), although the examples used in this report and in ART are only based on those in Theaetetus. The SOCRATES AS MIDWIFE view has been analyzed by Kittay (1989), and also by Holyoak & Thagard (1989) through the application to it of their ACME approach.

In section 4, ART distinguishes three different types of T → S with one type being subdivided into a positive and a negative version:

- T → S Transfer of Reasoning Queries, Focus or Expectation (ART section 4.1)
- T → S Transfer of (Un)certainty: Positive Influences (ART section 4.3)
- T → S Transfer of (Un)certainty: Negative Influences (ART section 4.4)
- Metaphorization of the Literal (ART section 4.5)

Detailed workings through of examples from each type are given in section 3 of the present report. In

[1] We prefer the term "metaphorical view" to Lakoff's term "conceptual metaphor." We take no stand here on the idea that metaphorical views are an important part of the very *concepts* tht people have for thinking the metaphors' target domains. Also we do not assume Lakoff's Invariance Hypothesis (Lakoff 1990), and do not use the notion of an image-schema.

section 1, the ATT-Meta system is described in sufficient detail to understand the type of rules and reasoning that are being used in section 3. In section 2, rules that represent the underlying source and target domain knowledge needed for the examples are given.

Finally, note that these user-defined rules are merely intended to show that ATT-Meta can model the types of Target to Source transfer described in the report. A notational variant of the rules described here has been supplied as user-defined rules to ATT-Meta, and the results and type of reasoning we report have been confirmed. What these rules are not intended to represent is the actual, far more complex, system of rules that would in reality be needed to handle metaphorical discourse involving the SOCRATES AS MIDWIFE metaphorical view.

1 The ATT-Meta System

1.1 Pretence Spaces

ATT-Meta is a reasoning engine or system that employs a process of *query-directed reasoning* (more commonly, but less clearly, called goal-directed reasoning). From an initial query –that is hypothesized to arise from (partial) processing of the co-text and other context– and using a set of user-defined *if-then* rules and facts about the source and target domains[2], the reasoning proceeds via repeated backwards applications of *if-then* rules until it is grounded in fact; a process known as backward chaining. In other words, a query Q is put to the system, which will hold if there is an *if-then* rule P → Q and a proposition P. Whether or not there is a rule P → Q can be determined by searching through the user defined rules. Whether or not the proposition P holds can be determined by treating it as a new query for the system to work on, and so on until the P matches a fact that the user has given to the system. If the search proves successful, then a *qualitative certainty level* such as *certainly, presumed, suggested* is then returned (c.f. section 1.2.2) along with possible instantiations of any variables existing in the query. A distinctive feature of ATT-Meta is that within the overall reasoning space, aspects of the reasoning may take place within nested reasoning spaces or cocoons, that is, within special computational environments; in these spaces, the beliefs pertaining to other agents may be simulated, or, the metaphorical utterance may be reasoned about as if it were believed to be literally true[3].

In this report, we shall deal solely with the metaphor-pretence-space. Those rules and facts supplied to ATT-Meta that might be characterised as belonging to the source domain would typically be used within this cocoon. Take, for example, the situation where a SOCRATES AS MIDWIFE metaphor has been used. When this view is in use, Socrates is viewed, counterfactually, as a midwife and mothers –who may not be mothers in reality– are viewed as being capable of giving birth to ideas. These views would hold within the metaphor-pretence-space. Then, the source domain knowledge and abilities of a typical midwife, such as their role in helping in a baby's birth, or in helping tend the baby in its first few days of life, could be taken to hold of Socrates and used for reasoning about Socrates and ideas within the metaphor-pretence-space.

Of course, what remains to be described is the issue of how the situation arises whereby Socrates is being reasoned about in a metaphor-pretence-cocoon using knowledge of the source domain. This is largely achieved via the use of *transfer rules*, which explicitly link a source domain predicate with a target domain predicate and state that the source domain predicate must hold within the metaphor-pretence space. Recall that ATT-Meta uses query directed reasoning. A target domain query (e.g. TQ) arising within reality may provide the input for a S → T (e.g. SQ → TQ[4]) transfer rule. This rule would state, in a manner we shall shortly examine, that a matching query (e.g. SQ) expressed in source domain terms also holds, and that it does so within the metaphor-pretence-cocoon. Therefore, via the process of query directed reasoning, Socrates will be reasoned about within the metaphor-pretence cocoon, using source domain rules about midwives. Given that the metaphor vehicle will have stated certain facts about Socrates as if he were a midwife –which we can model by requiring these facts to

[2] The terms *source-domain and target domain* are convenient for the purposes of describing a metaphor. However, ATT-Meta does not divide different types of knowledge into different domains. All it needs is an undifferentiated list of facts and rules.

[3] In both cases, knowledge of reality may appear to contradict the facts used or the conclusions reached within the cocoon. However, if this should occur, then the qualitative certainty level of information imported into the cocoon is downgraded.

[4] This is a simplification since the actual S → T rule will not involve SQ and TQ as such, but rather variables that will match the specific SQ and TQ.

hold only in the metaphor-pretence cocoon– then the reasoning within the pretence-cocoon should get grounded.

1.2 If-Then rules

1.2.1 Rules and Facts

The information manipulated by ATT-Meta consists of *if-then* rules and the facts, propositions, hypotheses or queries that result from or feed into the application of a rule. There are three main parts to a rule: the *if* part; the *then* part; and the *qualitative-certainty-level*.

The *if* part consists of a list, a conjunction, of conditions that must hold. These conditions are typically atomic, in that they have no internal connectives such as AND and OR, nor do they use logical quantification. Thus, the *if* part of a rule might look like the following, where the conjunction of atomic conditions is contained between square parentheses and the uppercase letters `P`, `I`, and `S` are variables:

(1) `[the_episode(gives_birth, P, I) the_episode(midwife, S)].`

The *then* part is usually an atomic condition or a negation of one. An example of such might be the following:

(2) `the_episode(midwife_for, S, P, I).`

Both *if* and *then* conditions are expressed in an episode or situation based logic (broadly similar in spirit to the logical scheme of Hobbs 1990). Since the reasons for using such a formalism are not important for this report, we will employ a different easier-to-read syntax, which doesn't refer to episodes and which expresses the conjunction of atomic conditions not in a list between parentheses, but by using the connective AND. Thus, (1) and (2) can be rewritten as follows:

(1') `gives_birth(P, I) AND midwife(S).`
(2') `midwife_for(S, P, I).`

In our syntax, we shall make the *if* and *then* explicit, so (1') and (2') can be combined into the following rule:

(3) `IF gives_birth(P, I) AND midwife(S) THEN midwife_for(S, P, I).`

This isn't quite the final version of the rule, since it does not yet contain a qualitative certainty level associated with the *then* clause. The certainty level is one of: *certainly, presumed, suggested*. The issue of how these certainty levels are used in reasoning will be addressed in the next section.

Suppose that we give rule (3) the level *suggested* so that informally we are stating that **if P has given birth to I and S is a midwife, then this suggests that S was the midwife for P**. Thus, rule (3) is rewritten as rule (4):

(4)
```
IF gives_birth(P, I)
AND midwife(S)
THEN {suggested} midwife_for(S, P).
```

In principle, any rule is applicable in any reasoning space[5]. However, it is possible to include a condition effectively demanding that the reasoning be in a particular space. For example, the following rule refers to both the metaphor-pretence-space and the reality-space.

(5)
```
IF currently-within-metaphor-pretence
AND in-reality-space: being_person(P)
AND in-reality-space: being_idea(I)
AND in-reality-space: produces(P, I)
```

[5] Rules with a certainty level of *certain* are downgraded to *presumed* when used in the metaphor-pretence-space.

```
THEN {presumed} gives_birth_to(P, I)
```

This rule is a reverse transfer rule (i.e., T → S). The first condition means that the rule can **only** succeed if it is being tried within the metaphor-pretence-space. The rule states that a hypothesis in metaphor-pretence-space `gives_birth_to(P,I)` will hold if a number of reality-space[6] conditions also hold; one in particular being `produces(P, I)`. It thus transfers information about the individuals that instantiate the variables `P` and `I` from the target domain to the source domain, and ensures that this only happens when the source domain is being reasoned about with the metaphor-pretence space, not the reality space.

Its converse, (6), illustrates the more normal situation of a source to target transfer (i.e., S → T).

(6)
```
IF being_person(P)
AND being_idea(I)
AND within-metaphor-pretence: gives_birth_to(P, I)
THEN {presumed} produces(P, I).
```

If this is applied within the reality space, then, if `being_person(P)`, and `being_idea(I)` are satisfied in that space, `produces(P, I)`, is supported in that space. The system looks down into the pretence space to see if `gives_birth_to(P, I)` is satisfied there.

Facts are just rules in which no *if* conditions are required to hold. Note that it is customary to give facts the certainty level of *certain*, although this is not required. The following two examples show a fact holding in any space, including metaphor-pretence, and one that will hold only in the metaphor-pretence.

(7)
```
{certain} being_person(electra)
```

(8)
```
within-metaphor-pretence {certain} midwife(socrates)
```

Fact (8) makes the counter-factual claim that **Socrates actually is a midwife**, and so allows the reasoning within the metaphor-pretence-space to be grounded. Recall that in a metaphor-pretence-space such claims are taken to be literally true.

1.2.2 Certainty levels

Given an initial hypothesis posed as a query, ATT-Meta returns *a qualitative certainty level* for the hypothesis after it has reasoned using the set of user-defined rules and facts. It can do this because all the facts and rules that ultimately support or do not support the query themselves have a qualitative certainty level associated with the *then* clause, and hypotheses, i.e. queries, established during reasoning have a certainty level assigned to them.

There are currently five certainty levels recognised by ATT-Meta for hypotheses: *certainly, presumed, suggested, possible,* and *certainly-not*. The meaning of *certainly* and its negation should be obvious. The certainty-level *presumed* might be read as 'by default'. The hypothesis is taken to hold until further evidence indicates otherwise. A hypothesis has the certainty level of *possible* if there is no evidence against it at the level of *certainly not*, but no evidence at all has been found to support it. The situation when an initial query is posed would be one such instance, and so initial queries are always given a certainty level of *possible*.

The certainty level of an *if-then* rule places an upper bound on the certainty level of the hypothesis corresponding to the *then* clause. The actual certainty level will usually be the minimum of the certainty levels of each of the hypotheses supporting the *if* clauses, and the certainty level of the rule. There are two related exceptions to this.

Firstly, the negation of a hypothesis is also investigated in case the evidence against it is at least as strong as that for it. When there is evidence to at least level *presumed* for both a hypothesis and its

[6]This is a simplification since pretence spaces may be nested, (c.f. Lee and Barnden (2001) for the use of nested spaces in handling mixed metaphors.). Consequently, the outer space might not be the reality space.

negation, then ATT-Meta's conflict-resolution mechanism will attempt to adjudicate between the two. This can result in one hypothesis winning, in which case the other's certainty is downgraded. Both hypotheses will have their certainty downgraded if after the attempt at conflict resolution, there is still no difference[7].

Secondly, if a hypothesis receives support from more than one *if-then* rule, then the maximum of the certainty values contributed by the different rules is used.

2 Domain Knowledge

In the remainder of this report the three types of T → S transfer will be illustrated. Short dialogues between **A** and **B** will be presented in which **A** will say something that includes a metaphor, and **B**'s response will indicate that (s)he has understood the metaphor in **A**'s statement and drawn an inference from it. In other words, *if **A**'s statement, then **B**'s statement follows*. We can think of this as a process of **B** posing to herself a query and looking and finding support for it in **A**'s statement. Thus, we can present **B**'s comment (in its logical form) as a query for ATT-Meta to work on.

In this section, we shall give rules to represent the underlying source and target domain knowledge assumed in the dialogues presented in the remainder of the report. Thus, we shall present some target domain rules, some source domain rules, and some transfer rules that map between the two. The dialogues will involve only the SOCRATES AS MIDWIFE metaphorical view. It must be emphasised that the rules presented here are intended to be merely illustrative. Consequently, they form a highly simplified and impoverished subset of the underlying knowledge about midwives (the source domain), Socrates and his acquaintances (the target domain) and the transfer rules mapping between the two domains that would actually be required.

The dialogues in section 3 of the report will introduce facts that will help ground these rules and so allow the reasoning to take place. These facts will be shown in the discussions following the different dialogues.

2.1 Target-Domain Rules

One of the main claims being made in the ATT-Meta project is that the importance of source-domain reasoning has been greatly under-appreciated in metaphor research. Thus, many of the examples that have been discussed within the project have emphasised the source domain rules. The examples discussed in this report are no exception and we require just one target domain rule.

(9)
```
IF helps_produce(S,P,I)
AND quarrel(S,P)
THEN {presumed} NOT(helps_develop(S,P,I)).
```

This rule represents the quite reasonable proposition that if someone (A) helps someone (B) produce an idea, but they subsequently quarrel, then it is unlikely that (A) will help (B) develop the idea. N.B., our implementation in this report does not handle time and change, so a rule such as this is artificially simplified.

2.2 Source-Domain Rules

The following two rules describe what it would mean for someone (S) to be a midwife for someone (P) at the birth of P's child (I). The first rule (10) states that if someone (e.g. Electra) gave birth to a child (her idea), and someone else (e.g. Socrates) had been mentioned as a midwife, then a tentative inference may be drawn that Socrates was the midwife for Electra at the birth of the idea. The certainty level of *suggested* is used because of the tentative nature of the inference.

(10)[8]
```
IF gives_birth(P,I)
```

[7] C.f. the discussion of reverse negative transfer of uncertainty in section 3.2.2, where an example of this mutual downgrading will be found.

[8] This is the rule 4 discussed earlier.

```
AND midwife(S)
THEN {suggested} midwife_for(S,P,I).
```

The next rule explicitly states that the individual who is the midwife (say Socrates) helped at the birth. Thus the `helps_give_birth(S,P,I)` predicate is used rather than the `gives_birth(P,I)`. The certainty-level is *presumed* rather than *certain* since Socrates might have been an off-duty midwife helping the actual midwife.

(11)
```
IF helps_give_birth(S,P,I)
AND midwife(S)
THEN {presumed} midwife_for(S,P,I).
```

The final rule in this series links the midwife with the birth by using a two place, rather than a one place, predicate `midwife(S,I)`, with the second argument matching the second argument of the first *if* clause. A certainty-level of *presumed* has been given, although arguably a level of *certain* might be justified. In general, we are hesitant in ascribing *certainty* to any rule, feeling that in the real world, most reasoning is uncertain. In the example that makes use of this rule (c.f. section 3.3), nothing rests on the exact certainty-level.

(12)
```
IF gives_birth(P,I)
AND midwife(S,I)
THEN {presumed} midwife_for(S,P,I).
```

The following two rules describe (some of) the inferences that can be made if it is known that someone (such as Socrates) is a midwife for another (such as Electra), possibly at the birth of the other's child (idea). Again, arguably, the certainty levels should be *certain* rather than the *presumed* that they have been given, but since the exact certainty level does not affect the examples that use these rules, we have remained cautious.

(13)
```
IF midwife_for(S,P,I)
THEN {presumed} helps_give_birth(S,P,I).
```

(14)
```
IF midwife_for(S, P, I)
THEN {presumed} helps_tend(S, P, I).
```

Rule (14) expresses the assumption that midwives not only help with the birth, but help the mother look after the baby and gives advice. Rule (15) states the rather obvious fact that if someone has helped someone else give birth, then that someone else must have given birth. We have stuck to our convention of giving a *presumed* certainty level. Rule (16) expresses the not unreasonable assumption that a baby that is receiving extra help is growing well.

(15)
```
IF helps_give_birth(S,P,I)
THEN {presumed} gives_birth(P,I).
```

(16)
```
IF helps_tend(S, P, I)
THEN {presumed} growing_well(I).
```

2.3 Transfer rules

Each of the following subsections contains four possible forms of each transfer rule[9], regardless of whether the dialogues in section 3 require all four forms or not. Thus, we have presented rules mapping from pretence to reality and from reality to pretence, together with their negative versions.

[9] There might appear to be a redundancy in the transfer rules that have been presented. Rule 17 maps between `gives_birth_to` and `produces`, and rule 21 maps between `helps_give_birth_to` and `helps_produce`. Rule 25 maps between `helps_tend` and `helps_develop`, and although there is no `tends` and `develops` mapping, it would be easy to imagine such a rule. Although the idea is not implemented in ATT-Meta, one might think of `helps_predicate` as a *View Neutral Mapping Adjunct (VNMA)* (see Barnden & Lee 2001) that can apply to all predicates involved in a transfer rule.

2.3.1 Correspondences Between Idea-Producing and Giving-Birth

(17)[10]
```
IF currently-within-metaphor-pretence
AND in-reality-space: being_person(P)
AND in-reality-space: being_idea(I)
AND in-reality-space: produces(P,I)
THEN {presumed} gives_birth_to(P,I).
```

(18)[11]
```
IF being_person(P)
AND being_idea(I)
AND within-metaphor-pretence: gives_birth_to(P, I)
THEN {presumed} produces(P, I).
```

(19)
```
IF currently-within-metaphor-pretence
AND in-reality-space: being_person(P)
AND in-reality-space: being_idea(I)
AND in-reality-space: NOT(produces(P,I))
THEN {presumed} NOT(gives_birth_to(P,I)).
```

(20)
```
IF being_person(P)
AND being_idea(I)
AND within-metaphor-pretence: NOT(gives_birth_to(P, I))
THEN {presumed} NOT(produces(P, I)).
```

2.3.2 Correspondences Between Helping to Produce Ideas and Helping to Give-Birth

(21)
```
IF currently-within-metaphor-pretence
AND in-reality-space: being_person(S)
AND in-reality-space: being_person(P)
AND in-reality-space: being_idea(I)
AND in-reality-space: helps_produce(S,P,I)
THEN {presumed} helps_give_birth(S,P,I).
```

(22)
```
IF being_person(S)
AND being_person(P)
AND being_idea(I)
AND within-metaphor-pretence: helps_give_birth(S,P,I)
THEN {presumed} helps_produce(S,P,I).
```

(23)
```
IF currently-within-metaphor-pretence
AND in-reality-space: being_person(S)
AND in-reality-space: being_person(P)
AND in-reality-space: being_idea(I)
AND in-reality-space: NOT(helps_produce(S,P,I))
THEN {presumed} NOT(helps_give_birth(S,P,I)).
```

(24)
```
IF being_person(S)
AND being_person(P)
AND being_idea(I)
AND within-metaphor-pretence: NOT(helps_give_birth(S,P,I))
THEN {presumed} NOT(helps_produce(S,P,I)).
```

2.3.3 Correspondences Between Helping to Develop Ideas and Helping to Tend Infants

(25)
```
IF being_person(S)
AND being_person(P)
AND being_idea(I)
```

[10] This is rule (5) discussed earlier.

[11] This is rule (6) discussed earlier.

```
      AND within-metaphor-pretence: helps_tend(S,P,I)
      THEN {presumed} helps_develop(S,P,I).

(26)  IF currently-within-metaphor-pretence
      AND in-reality-space: being_person(S)
      AND in-reality-space: being_person(P)
      AND in-reality-space: being_idea(I)
      AND in-reality-space: helps_develop(S,P,I)
      THEN {presumed} helps_tend(S,P,I).

(27)  IF being_person(S)
      AND being_person(P)
      AND being_idea(I)
      AND within-metaphor-pretence: NOT(helps_tend(S,P,I))
      THEN {presumed} NOT(helps_develop(S,P,I)).

(28)  IF currently-within-metaphor-pretence
      AND in-reality-space: being_person(S)
      AND in-reality-space: being_person(P)
      AND in-reality-space: being_idea(I)
      AND in-reality-space: NOT(helps_develop(S,P,I))
      THEN {presumed} NOT(helps_tend(S,P,I)).
```

2.3.4 Correspondences Between an Idea Developing Well and an Infant Growing Well

```
(29)  IF being_idea(I)
      AND within_metaphor_pretence: growing_well(I)
      THEN {presumed} developing_well(I).

(30)  IF currently-within-metaphor-pretence
      AND in-reality-space: being_idea(I)
      AND in-reality-space: developing_well(I)
      THEN {presumed} growing_well(I).

(31)  IF being_idea(I)
      AND within_metaphor_pretence: NOT(growing_well(I))
      THEN {presumed} NOT(developing_well(I)).

(32)  IF currently-within-metaphor-pretence
      AND in-reality-space: being_idea(I)
      AND in-reality-space: NOT(developing_well(I))
      THEN {presumed} NOT(growing_well(I)).
```

2.4 Facts about all the dialogue Situations

```
(33)  {certain} being_person(socrates).

(34)  {certain} being_person(electra).

(35)  {certain} being_idea(idea).
```

3 Target to Source Transfer

In this section, the three types of target to source transfer discussed in section 4 of ART are illustrated.

[12] This is rule (7) discussed earlier.

3.1 Transfer of Reasoning Queries

The first case of target to source transfer can be illustrated very simply within ATT-Meta, since the operation of ATT-Meta proceeds in exactly this query-directed (or goal-directed) fashion. That is, the situation we wish to model is one in which the discourse prior to a metaphorical utterance poses an issue which we frame as a query and the metaphorical utterance is assumed to provide (part of) the answer. Since the answer is couched in target-domain terms, the query needs to be converted to one that is also couched in target-domain terms[13]. This is what transfer rules do in ATT-Meta. Once a top-level query has been provided, ATT-Meta will attempt to match it with the *then* conditions(s) of an *if-then* rule. Unless the matching rule had been a fact with no *if* conditions, the *if* conditions would then represent subqueries that ATT-Meta would attempt to match with the *then* conditions of further rules or with facts, and so on recursively.

Suppose that a short dialogue had gone something along the following lines:

(A) Electra mentioned an interesting idea to Orestes. Once again, Socrates had been her midwife.
(B) So, Electra must have produced the idea.

A's comments introduce the following fact amongst others, which will be used by B, together with some of the general rules and facts from section 2, to make inferences:

(36)
```
within-metaphor-pretence {certain}
      midwife_for(socrates, electra, idea).
```

Let us make the fairly natural assumption that B, upon hearing that Electra mentioned an interesting idea to Orestes, would wonder whether Electra was the originator of the idea or was merely repeating something she had heard or read. In other words, B might pose a query that could be expressed in the formal language of ATT-Meta (as simplified in this report) as 37:

(37)
```
produce(electra, idea).
```

B's query is answered, in A's second sentence, with a metaphor. B has clearly taken this sentence as a confirmation of the query since (s)he concludes "So, Electra must have produced the idea". In conclusion, the question and response must have been integrated in order for the question to be answered. Or, in other words, the *target-query* must have been converted into a query within the *metaphor-pretence* domain.

Such a process would proceed in the following manner. The *target query*, (i.e 37) matches the *then* clause in the transfer rule (18), with the variables `P` and `I` instantiated as `electra`, and `idea`, respectively.

By a standard process of back-chaining, this then sets up three subqueries. The first two belong to the *reality-domain*.

(38)
```
being_person(electra)
```
(39)
```
being_idea(idea)
```

Rule (34) supports the first of these target level queries, and rule (35) the second.

The third subquery represents a *metaphor-pretence-level* query.

(40)
```
within-metaphor-pretence: gives_birth_to(electra, idea).
```

Thus, the system tries to establish the query (40) within the metaphor-pretence. Since this new query is expressed in source domain terms, the query has been transferred from the target to the source domain. Notice carefully that the transfer is mediated by rule (18), which expresses a mapping relationship from

[13] Currently, ATT-Meta cannot itself frame a reasoning initiating question from discourse. Instead, it must rely upon the user of the system to frame such a query.

source to target. Thus we have an instance of reverse transfer, although the mapping is source to target[14].

Let us assume that the hearer and poser of the question are familiar with the source domain rules (15) and (13). The former states informally that **if someone has helped you to give birth to a baby, then you have given birth to a baby**. The latter states that **if someone is your baby's midwife, then they have helped you to give birth to your baby**.

The *then* clause of rule (15) matches query (40), i.e., the *if* clause of (18), with the variables `P` and `I` instantiated as `electra` and `idea` respectively. As a consequence of this match, the *if* clause of (15) is posed as a new query:

(41) `helps_give_birth_to(S,electra,idea).`

This matches the *then* clause of rule (13), which means that the *if* clause is now posed as a new query. It gets its answer from the second sentence in the discourse, which is represented by fact (36). Thus the target domain query (37) is grounded by the source domain fact (36), which was introduced into the discourse by A.

Note that in this example, only a little reasoning has been done within the *metaphor-pretence-space*. Furthermore, the transfer of certainty levels has been straightforward. The next examples we shall discuss require more reasoning, part of which will involve conflicts between rules with different levels of certainty.

3.2 Reverse Transfer of (Un)certainty

As described in section 1.2.2, given a query, which may be the initial query, but may also be any other query created during the process of reasoning, ATT-Meta will return a certainty level for that query. This will have been derived from the interaction of the certainty levels of the rules and facts used to support the query. When part of that support rests upon a S $\rightarrow$ T transfer of a source-domain proposition, the certainty level of that proposition can affect the certainty level of the target-domain that will have been created by the transfer. However, in ART it is argued that the reverse may also occur: additional support for the newly created target-domain proposition coming via other chains of reasoning should be transferred back to the source-domain proposition. In turn this raised or lowered certainty level of the source-domain proposition may influence the certainty-level of any further proposition that is derived from the source-domain proposition. Notice carefully that this process shows the reverse transfer of a certainty level, even though it was the target proposition that was created from the source and not the other way around. We shall now deal with a positive and a negative version of this reverse-transfer.

3.2.1 Reverse Transfer of Positive Influences

We shall first illustrate how ATT-Meta models the transfer of positive influences on certainty from target to source.

Consider the following discourse:

(A) Electra gave birth to an extremely interesting idea, the other day. I must say, Socrates really is an excellent midwife. It was he, of course, who helped her come up with the idea in the first place.
(B) So, he must have helped her develop it too.

A's comments introduce the following facts amongst others, which will be used by B, together with some of the general rules and facts from section 2, to make inferences:

(42) `within-metaphor-pretence {certain} gives_birth(electra, idea)`

[14] ATT-Meta can contain target-to-source rules such as (17), and these will be used in the following sections. However, these are not the issue in this subsection.

(43)[15] `within-metaphor-pretence {certain} midwife(socrates).`
(44) `{certain} helps_produce(socrates, electra, idea).`

B's response to A, in particular, her/his use of the word 'so', suggests that B has at this point posed the question to her/himself as to whether Socrates helped Electra develop the idea, and found an answer for it in A's statements. In other words, (s)he has posed the target domain query:

(45) `helps_develop(socrates, electra, idea)`

and found a chain of reasoning that supports it. This, as we shall show, depends on the knowledge that if someone is a midwife for a mother (such as Socrates for Electra), then not only will (s)he help with the birth, but will also help tend the child in its first few weeks of life. A transfer rule will then link helps_tend and helps_develop.

What is interesting about this dialogue is that the source domain aspects of A's comments, which are found in the first two sentences, do not explicitly state that Socrates was Electra's midwife for the birth. Consequently, we shall see that at one stage of reasoning the certainty level for the hypothesis that Socrates is the midwife for Electra is only *suggested*. However, we shall demonstrate that if A's third sentence, in which Socrates and Electra and the idea are explicitly linked, is metaphorized[16], then a boost in certainty can be given to this hypothesis, and ultimately to the target-level query. This justifies B's confident response that Socrates helped Electra develop the idea.

The query (45) is expressed in terms of the target domain and needs converting into source domain terms in order to find support from the claims about Socrates being a midwife. Transfer rule (25) (**if someone S helps someone else P to tend a child I, then S helps P develop I)** achieves this.

Once the variables of (25) are instantiated, the first three clauses all correspond to facts: (33), (34), and (35), which leaves the fourth clause providing a source domain query within the pretence:

(46) `within-metaphor-pretence: helps_tend(socrates, electra, idea).`

A suitable match is provided by the *then* clause of rule (14), i.e. **if someone S were a midwife for someone else P and her baby I, then S would help P tend the baby I after its birth**. The *if* clause, once instantiated with Socrates and Electra and Electra's idea, requires Socrates to have been the midwife for Electra and her idea.

(47) `midwife_for(socrates, electra, idea).`

This new query does not correspond to any fact stated in the discourse. Instead, we just have the claim that **Socrates was a midwife** and that **Electra gave birth to an idea**, which can be represented in ATT-Meta as the facts (43) and (37). However, these two facts match the *if* clauses of rule (10) with the *then* clause matching (47). Note that the certainty level of this rule is not the *presumed* of the other rules that have been discussed so far, but the lower one of *suggested*.

Consequently, this chain of reasoning allows ATT-Meta to conclude with a certainty level of *suggested* within the pretence that **Socrates was the midwife for Electra**. Since this proposition (i.e., (47)) matches the *if* clause of rule (14), ATT-Meta must conclude that **Socrates helped Electra tend the idea** with certainty only *suggested*, unless the proposition that corresponds to the *if* condition of rule (14), can be supported to the level of *presumed* by an alternative chain of reasoning.

Without this alternative chain of reasoning returning a higher certainty-level, then the *suggested* certainty level pertaining to the source level query, (46), will in turn be passed to the initial target-level query, (45), in reality space, which does not seem to fully justify B's confident response to A.

However, there are aspects of the discourse that have not yet been considered. In particular we have not yet examined the role of fact (44), i.e., the target-domain statement that **Socrates helped Electra come**

[15] This is rule (8) discussed earlier

[16] In ART, reverse transfer of positive certainty influences is treated separately from metaphorization (see ART section 4.4), but in ATT-Meta the former becomes a special case of the latter.

up with the idea.

We can assume that the participants in the discourse have knowledge of transfer rule (21), which establishes a correspondence between **helping to produce ideas** and **helping give birth**. We can also assume knowledge of the source-domain rule (11), which states, informally, that **if someone helps a mother give birth to an idea and that someone is a midwife, then (s)he is the midwife for the mother**. Unlike the case in rule (10), where the certainty level was *suggested*, the certainty level in (11) is *presumed*.

The presence of rules (21) and (11) mean that there are now two arguments in support of the proposition that **Socrates was a midwife for Electra**, i.e. (47). One originating in the source domain with a certainty level of *suggested*, and one originating in the target domain with a certainty level of *presumed*. Since, the certainty level of a proposition is the maximum of the certainty levels of the arguments supporting the proposition, the certainty-level of (47) will be *presumed*. This certainty level is passed to the source query (46), (i.e., **Socrates helped Electra tend the idea**), and this in turn is passed to the initial target query (45) (i.e. **Socrates helped Electra develop her idea**).

In summary, (21) enables a reverse transfer of a positive certainty influence, and the result of this reverse transfer affects reasoning within the source domain and ultimately causes a forward transfer of positive influence to the initial target query.

3.2.2 Reverse transfer of negative influences

Consider the following short discourse:

(A) Electra gave birth to an extremely interesting idea, the other day. Once again, Socrates was her midwife, but they soon quarrelled.
(B) So, we can't be sure it is developing well.

A's comments introduce the following facts amongst others, which will be used by B, together with some of the general rules and facts from section 2, to make inferences:

(48) `{certain} quarrelled(socrates, electra).`
(49) `within-metaphor-pretence {certain}`
`midwife_for(socrates,electra, idea).`[17]

On the basis of what A has said, it would appear that B has mentally posed the question as to whether the idea is developing well and concluded that there are no strong arguments either way.

(50) `developing_well(idea).`

Transfer rule (29) states informally that **if, within the metaphor pretence, something is growing well, then it is also developing well**. Thus, query (51) is posed within the metaphor-pretence cocoon.

(51) `growing_well(idea).`

The source domain rule (16), states that **if someone helps someone tend something, then it will grow well**, and as we saw in the previous section, **if someone were someone's midwife for a child, then they would tend the child after its birth**, (rule 14). However, unlike the situation in the previous section, the discourse here makes it clear that Socrates was the midwife for Electra. Consequently, the *if* condition of rule (14) can be grounded in a given fact, (49).

Thus, there is a chain of arguments supporting the initial query (50). Since all the rules used in support have a certainty-level of *presumed*, B's response would appear to show undue caution unless there were reasons to assume that Socrates might not be helping Electra tend the idea.

If there are arguments both for and against a proposition, then the conflict-resolution mechanism

[17] In section 3.1, this fact was numbered (36), and was taken to underlie A's slightly different comments.

discussed in section 1.2.2 will attempt to arbitrate. The certainty-level of (52), the positive form of the argument, is *presumed*. If the certainty-level of (53), the negative form of the argument, is also *presumed*, and one proposition does not rely upon more specific arguments than the other, then, both propositions will have their certainty-levels lowered to *suggested*.

(52) `helps_tend(socrates, electra, idea).`
(53) `NOT(helps_tend(socrates, electra, idea)).`

Suppose that this were the case. Then the certainty-level passed to (51) will be *suggested*, and this level will be passed in turn to (50) i.e. the initial query. This certainty-level is appropriate for B's response. So what are the arguments supporting (53)?[18]

Transfer rule (28) (**if someone S doesn't help someone P develop I, then S will not help P tend I**) mapping from reality to the metaphor-pretence supports the proposition, and raises in turn the target domain query:

(54) `NOT(helps_develop(socrates, electra, idea)).`

This can be answered by rule (9) in the target domain which states informally that **if someone S helps someone else P produce an idea I, but the two subsequently quarrel, then it is likely that the former will not help the latter develop the idea.**

Fact (48) satisfies the second *if* clause of (9), but an instantiation of the first *if* clause is not found in the dialogue. So instead, the following query is posed:

(55) `helps_produce(socrates, electra, idea).`

This is a target domain query, and there is no target-domain information in our discourse that would support it. However, the use of transfer rule (22), (**if S helps P give birth to I, then S helps P produce I**) will result in the following source-domain query being posed within the metaphor-pretence cocoon.

(56) `helps_give_birth(socrates, electra, idea).`

This will match the *then* clause of rule (13) (**if someone is the midwife for someone else's baby, then (s)he will have helped that someone else give birth to the baby**), and the *if* clause will match fact (49) from A's remarks.

Finally, note that the certainty-level has been *presumed* throughout this chain of reasoning supporting query (53).

In summary, both proposition (52) and proposition (53) are supported with a certainty-level of *presumed*. The conflict-resolution mechanism is unable to find a winner and, as a consequence, the certainty-level of both is downgraded to *suggested*. This level is ultimately passed to the initial query, (50). The latter proposition was supported because transfer rule (28) enabled negative information in the target domain to have an effect in the source domain. But, note that the situation is even more complicated, since the negative influences did not just come from the target domain. In order to have an effect, proposition (48) needed to be coupled with hypothesis (55), which stemmed via a transfer rule from source domain hypothesis (56). Crucially, this hypothesis was not used directly within the source domain. This passing of information backwards and forwards between source and target domains is, we would claim, an important aspect of the use of metaphors.

3.3 Metaphorization of the Literal

All of the examples discussed in this report illustrate to some extent the metaphorizing of the literal. When a hypothesis in the target domain is converted via a transfer rule to a source domain hypothesis, then it has undergone metaphorization. To give a further illustration, let us make a slight change to the

[18] This query will be posed by ATT-Meta because whenever a query is investigated (e.g. `helps_tend(socrates, electra, idea)`), the negative version is simultaneously investigated in order to find whether its certainty level is higher or equal, c.f. section 1.2.2.

discourse discussed in section 3.2.1, Reverse Transfer of Positive Influences.

(A) Electra produced an extremely interesting idea, the other day. Of course, Socrates was its midwife.
(B) So, he must have helped her develop it.

Unlike the situation in the previous passage, Socrates is explicitly mentioned as being the idea's midwife here, or in ATT-Meta terms:

(57) `within-metaphor-pretence {certain} midwife(socrates, idea).`

This fact matches the second *if* clause of the source-domain rule (12), which states informally that **if someone P gives birth to an idea I, and someone S is the midwife for I, then S is Ps midwife for I**. Consequently, if proposition (58), i.e., the first *if* clause, were known or could be inferred, then it could be concluded within the metaphor-pretence that proposition (59) holds.

(58) `gives_birth(electra, idea).`
(59) `midwife(socrates, electra idea).`

This, in turn, would allow, via the use of rule (14) in the metaphor pretence (**if S is Ps midwife for I, then S helps P tend I**), the following source domain conclusion to be drawn:

(60) `helps_tend(socrates, electra, idea).`

The use of transfer rule (25) transfers this proposition to the reality space (**if S helps P tend I, then S helps P develop I**) and this target domain conclusion justifies B's response to A.

However, this chain of reasoning depends upon there being a proposition in the metaphor pretence, and unlike the situation in the other dialogues that have been discussed, A's comments do not include this source domain fact. A, instead, makes a target domain claim that **Electra produced an extremely interesting idea**, or in ATT_Meta terms:

(61) `produces(electra, idea).`

If this statement could be interpreted in terms of the Socrates As Midwife metaphor, as being about giving birth, then the chain of arguments would be supported. Transfer rule (17) (**if P produces I, then P gives birth to I**) achieves this metaphorization[19].

Acknowledgment

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

References

Barnden, J.A., Glasbey, S.R., Lee, M.G. & Wallington, A.M. (2002). Asymmetry and Reverse Transfers in Metaphor. Technical Report CSRP-02-9, School of Computer Science, The University of Birmingham, UK.

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, UK.

[19] It must be stressed that the rules assumed in this report, together with the dialogues are merely intended to illustrate how ATT-Meta could deal with the three types of reverse transfer. Thus, we are not claiming that all readers of this dialogue would metaphorize the proposition that Electra produced an extremely interesting idea. It is possible that some would convert the metaphorical term *midwife* to the target domain term *teacher*.

Hobbs, J.R. (1990). *Literature and cognition*. CSLI Lecture Notes, No. 21, Stanford University.

Holyoak, K.J. & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science, 13* (3), pp.295-355.

Kittay, E.F. (1989). *Metaphor: its cognitive force and linguistic structure*. (Paperback ed.) Oxford, UK: Clarendon Press.

Lakoff, G. (1990). The Invariance Hypothesis: is abstract reason based on image-schemas? *Cognitive Linguistics, 1* (1), pp. 39-74.

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.

Plato (1973). *Theaetetus* translated (from the Greek) with notes by John McDowell. Oxford: Clarendon Press.