

Unparalleled Creativity in Metaphor

John A. Barnden

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

Abstract

We address some issues concerning the relationship between metaphor and creativity, arising from an AI project (ATT-Meta) that has developed a partially-implemented theory of metaphor understanding. The central claim is that, while metaphor is plausibly based partly on structured mappings (analogies) between a target subject matter (the topic actually being addressed) and some source subject matter, creative metaphorical phraseology very often does not require construction of new mapping links to handle source aspects used by the phraseology but not handled by existing mappings. Thus, such phraseology often rests on substantial *non-parallelism* (non-analogy) between source and target. Rather, the unmapped source aspects serve only to indirectly control what information is transferred between source and target by already-known mapping links; and the phraseology, while itself creative, is not causing a grandly creative new look at the target. However, the approach explains relatively easily the nature of much creative metaphor and how it can be understood, and has some implications for creativity more generally.

Introduction

Much of the exploration of metaphor in fields such as Cognitive Linguistics and natural language Pragmatics takes metaphor to rest on complex mappings between a target subject matter and a source subject matter (see, e.g., Lakoff 1993, Kövecses 2002). These mappings constitute a structured analogy between some sector of the source subject matter and some sector of the target subject matter. To put it another way, some aspects of the source, generally including some structural aspects, are deemed to have parallels in the target. It is typical for the whole of the meaning of a metaphorical expression to be explained by means of these parallels.

While much of the emphasis in such accounts is on parallels that the understander is already familiar with, a prominent line of work on Psychology and AI on detailed processing models of metaphor has been on discovering the source/target parallelism (analogy) from scratch. Thus, psychological or AI models of analogy construction such as SME (Falkenhainer, Forbus and Gentner 1989) and ACME (Holyoak, Novick and Melz 1994) have been applied to metaphor. The AI approaches to metaphor of Wilks (1978)

and Fass (1997) have also been based on mapping creation, as has one part of Hobbs's (1990) approach.

The above approaches taken by themselves give the impression that the task of metaphor understanding is to use the fullest analogy obtainable between the target and the parts of the source that the metaphorical expression uses, whether the analogy is known in advance, discovered from scratch, or partly known in advance and partly discovered from scratch. This stance has a particular consequence, concerning aspects of the source subject-matter that are in a metaphorical expression but that are not paralleled in the target (via a metaphorical mapping that the understander already knows). Presumably, the understander would have to extend the known analogy by creating a new set of mapping links to handle the unparalleled source aspects (and also to map the connections they have to other source aspects that are mapped, so as to get as much structural parallelism between source and target as possible). This may in turn necessitate the stipulation of the existence of new target aspects, not already present in the understander's conception of the target, so that the unmapped source aspects have something to map to.

In contrast, the author's own work on metaphor, in an AI project called ATT-Meta (Barnden and Lee 2001, Barnden 2001, 2006, Barnden, Glasbey et al. 2004), is an exponent of a camp that, while giving some strong role to parallelism in metaphor, nevertheless tends to refrain from assuming that unparalleled source aspects should be given freshly-created parallels. Instead, the camp seeks to rely on possibly-extensive inference in the terms of the source subject-matter to link the unparalleled aspects to source aspects that do have a known mapping. This approach allows complex, controlled transfer of information from source to target without the overhead of establishing an extended analogy between source and target or worrying about whether an extended analogy even exists. On occasion, an unparalleled source aspect may need to be given a target correspondent (or a given understander may in fact give it such a correspondent even if it is not necessary to do so), but this is the exception not the rule.

Other approaches aside from ATT-Meta in this camp include another part of Hobbs's account, Narayanan's account (Hobbs 1990, Narayanan 1999), and to some extent at least the partly Narayanan-inspired neural metaphor theory devel-

oped by Lakoff and co-workers (see Afterword in Lakoff and Johnson 2003). Also, Langlotz (2006) provides a detailed and extensive treatment of (non-)parallelism in addressing issues of systematicity, transparency and motivatedness of metaphorical idiom. The camp in general allows substantial non-parallelism between the source aspects used by a metaphorical expression and the target aspects addressed. However, the ATT-Meta approach has pushed the idea further and more explicitly than other approaches have. In particular it has emphasized that, often, not only are the un-mapped source aspects *not* mapped but that they may not *able* to be mapped in any non-vacuous and genuinely useful way.

Not all theories of metaphor are based on parallelism at all, at least ostensibly. Two main exceptions are Relevance Theory (Sperber and Wilson 1995) and the categorization model of Glucksberg and Keysar (1990). But, while to some extent sketching processes that lead from utterances and other information (e.g., context and world knowledge) to metaphorical meanings, they do not explicitly engage with the details of meaning extraction such as those arising in our treatment of examples below where we claim non-parallelism.

The advocacy of non-parallelism in ATT-Meta and other approaches in the surrounding camp is restricted to cases where the metaphorical expression is indeed partly based on some already-known analogy. In some cases of metaphor the metaphorical utterance does not rest on any known analogy. This arises most clearly in many instances of novel “image metaphor,” where some novel visual (or other perceptual) similarity is discerned between source and target. We do assume that in thoroughly novel metaphor, whether image-based or not, the understander will often need to create an analogy for him/her/itself.

Putting aside completely novel metaphor, the ATT-Meta approach allows still-highly-creative instances of metaphor to involve non-parallelism. The unparalleled source aspects serve, via chains of inference linking them to mapped source aspects, to control what is transferred to the target. In that this transfer is by already known mappings, it is fair to say that no grand, new, creative view of the target is generated. However, the effect is that the richness of the source subject matter can be used very flexibly to control precisely what is transferred, and to lead to complex, subtle information about the target being conveyed. Such information may be difficult or practically impossible to convey without the metaphor (this is the often-discussed unparaphrasability of metaphor: see, e.g., Stern 2000).

In the next section we will examine some examples of rich non-parallelism in metaphor. The section after that describes the ATT-Meta theory and system through one of those examples. Then there is a discussion and conclusion section that further relates the preceding to creativity issues.

We will use the word “parallelism” rather than the more technical word “isomorphism,” or even “analogy,” because “isomorphism” implies a strict one-to-one correspondence between items in the source scenario and items in the target scenario, and analogy theorists often hold that analogy rests on isomorphism. We wish not to prejudge the question of

whether looser, messier forms of parallelism are sometimes needed.

Examples of Non-Parallelism

Nylon Strings

Consider the following spoken comment on a TV programme (*Newsnight*, BBC2 channel, U.K., 3 July 2007; plausible punctuation added):

(1) I don't think strings are attached. If there are any they're made of nylon—I can't see them.

The speaker was an African politician being interviewed about a new investment by China. In the first sentence, he uses stock metaphorical phraseology concerning strings being attached. We take the “strings” to refer to constraints affecting the politician's country's future behaviour with regard to the investment (cf. the meaning for “no strings attached” given in the *Oxford Dictionary of Idioms*, 2005, namely that no special restrictions or conditions apply to an opportunity or offer). In the second sentence, the politician then elaborates on the standard phraseology. We take him there to be saying, in source terms, that he can't “see” any “strings” because if they exist they're made of nylon; and thus to be saying in target terms that he is not aware of any constraints, and that if any exist they must be somehow unnoticeable.

How much parallelism (structured analogy) should the understander seek between the source scenario that is painted and the target scenario that is actually being portrayed? In the source scenario there are strings, attachment of them (to the African parties involved in the business investment in question, metaphorically cast as physical objects that are somehow manipulable by strings), the substance nylon, the physical property of being made of a substance, and the ability of the speaker to physically see the strings. Which of these aspects of the source scenario correspond to anything at all in the target scenario? The strings have a parallel (the constraints), but what about the nylon, for instance?

Before going on we must note a complication. It could be argued that the interpretation of “I can't see them” could be separate from the interpretation of the prior parts of the passage, and be based on interpreting “them” directly as referring to something in the target scenario corresponding to the physical strings, rather than to the physical strings themselves. If this is how understanding proceeds then it is furthermore possible that the understander can appeal directly to a non-physical lexical sense of “see”—e.g., a sense of abstractly being mentally aware of something; in which case that non-physical scenario needs to be entertained in understanding “I can't see them.” We will put these possibilities aside, as they would only force small amendments to our discussion, and in view of the patent coherence between the physical translucency or invisibility of nylon and the physical process of seeing.

We do assume that the (possible) strings and the (possible) attachment of them in the source do have parallels in the target: the strings correspond to possible behaviour-affecting constraints that the Chinese investors might have sought to place as a condition of the investment, and the attachment is

that (counterfactual) placing and behaviour-affecting itself. So what remains is the nylon, the being-physically-made-of, and the being-able-to-physically-see. In fact, in the following we break the last item down into being-able and the physical seeing itself. The main claim about the example is that it mainly seeks to convey that the speaker is

not able to notice any extra constraints affecting the behaviour of the African side.

Here, the (abstract) noticing corresponds to (i.e., is the parallel of) seeing, and the not-being-able associated with the noticing in the target scenario corresponds to the not-being-able associated with the seeing in the source scenario. One can also interpret the passage as conveying that

the speaker is allowing for the possibility that there are constraints, but if so then they are of a nature such that it is intrinsically difficult for anyone to notice them .

That is, it is not just inattention, stupidity, or whatever on the part of the speaker that he can't "see" them.

How is the above difficulty in the target scenario arrived at by the understander? Presumably from difficulty of physically seeing nylon strings. And where does that difficulty come from? The translucency of nylon, combined possibly with the thinness of nylon strings. The crucial carry-over here is that of difficulty. Notice carefully that this does not of itself imply that the translucency needs to be mapped into the target scenario. Even less does it imply that being-made-of-nylon itself needs to be mapped (similarly for being thin) and, even less again, that being-made-of and nylon need separately to be mapped.

Of course, if one asked the question of why the (possible) constraints in the target scenario are difficult to notice, one might seek some target parallel of translucency. But certainly there was nothing in the actual context of the example as heard that could provide any specific target parallel for translucency. The best that can be suggested is that *some property or other* of the constraints makes them difficult to notice. But the understander could work out that totally vacuous and useless piece of information without any metaphorical utterance or any other information being at hand.

So, we have arrived at the suggestion that a rich understanding of the passage can be obtained on the basis merely of the following parallels:

strings — behaviour-affecting constraints
attachment — the application of those constraints
physical seeing — abstract noticing
inability — inability
difficulty — difficulty
possibility, counterfactuality, etc. — possibility, counterfactuality, etc. (respectively)

A point we will come back to later is that only the first three of these are specific to the the nature of the source and target at hand.

Conversational Cat Flaps

Barnden et al. (2004) extensively analysed the following example:

(2) I tried not to run down Phil too much - I felt bad enough as it was, what with screwing his girlfriend and all. But it became unavoidable, because when Jackie expressed doubts about him, I had to nurture those doubts as if they were tiny, sickly kittens, until eventually they became sturdy, healthy grievances, with their own cat-flaps which allowed them to wander in and out of our conversation at will. [Hornby, *High Fidelity*, Penguin, 1995, p. 12]

The claim is that it would be a mistake (a waste of time) to try to find parallels for the following source-scenario elements (or possibly others):

the cat-ness of the doubts and grievances
their sickliness
the cat-flaps
the deliberateness of the wandering (cf. the "at will").

The most obvious case here is the cat-flaps. The point about these is that they support the point (which is also explicitly stated) that the cats can wander in and out at will, this wandering corresponding to the grievances appearing intermittently in the conversation. It would seem pointless to try to work out some specific aspect of the conversation or some other part of the target scenario that corresponds to the cat-flaps. Also, the deliberateness of the cats' wandering has no parallel with any feature of the grievances corresponding to the cats. Rather, the deliberateness of the cats' wandering suggests, *within the source scenario*, that it is not the speaker who is bringing the cats in and out: that is, the cats' movements are non-deliberate on the part of the speaker. It is *his non-deliberateness* that maps: it maps to the (alleged) non-deliberateness on his part with which the grievances appear in and disappear from the conversation.

The main point of the sickliness (and tininess), arguably, is to buttress the metaphorical notion of "nurturing" the kittens. Clearly this nurturing itself does have some parallel in the target scenario. Now, if the kittens are sickly and tiny then nurturing is especially natural and morally good. Perhaps therefore the point of the sickliness and tininess is to portray the speaker's action in converting doubts to grievances as natural and morally good. More securely, it can be inferred in the source scenario that the kittens were helpless and were likely to die if not nurtured. From this, we might suppose that in the target scenario the doubts were likely to disappear unless the speaker acted. So, the tininess and sickliness certainly lead to source-scenario conclusions that map to the target in some way, generating important information there, but that does not mean that they themselves, or the steps connecting them to those conclusions, map to the target.

Finally for the above list, the cat-ness is totally irrelevant except in that it underlies familiar domestic scenarios of nurturing kittens, cats using cat-flaps, and so on. So it contributes to the development of a coherent source scenario, but there is no reason at all to suppose that some specific properties of cats (aside from the being-nurtured, the wandering, etc.) are to be mapped to the target.

Other Examples

An example we have treated in various papers (Barnden 2001, Barnden, Glasbey, *et al.* 2004, Barnden and Lee 2001) and implemented using the ATT-Meta system is the following:

- (3) In the far reaches of her mind, Anne knew Kyle [her husband] was having an affair, but ‘to acknowledge the betrayal would mean I’d have to take a stand. I’d never be able to go back to what I was familiar with,’ she says. Not until eight months had passed and she finally checked the phone bill did Anne confront the reality of her husband’s deception. [In L. Gross, “Facing up to the Dreadful Dangers of Denial,” *Cosmopolitan* 216(3), USA ed., March 1994.]

Here we claim that the point of the “far reaches” is to generate the *within-source* conclusion that Anne’s conscious self (metaphorically conceived as a person located in the middle of Anne’s mind-space) has great difficulty in physically operating upon the idea that Kyle was having an affair (this idea being metaphorically cast as a physical object physically located in the far reaches of the mind-space). There is no need to, and it would probably be very difficult to, find some component or aspect of Anne’s mind that corresponds to (parallels) the far reaches themselves. Again, we could postulate some such component, and say that the affair-thought was in some sense “in” it, but what would be the point? The message of the utterance is the fact that Anne is in a mental state of having difficulty in consciously involving the affair-thought in her thinking. And this difficulty comes from the difficulty of, in metaphorical terms, physically operating upon the idea. The “far reaches” are just a within-source tool for conveying this physical difficulty.

Barnden (2006) discusses the degree of non-parallelism in fifteen further examples taken from real discourse. One of those examples is:

- (4) This all means that general managers have cricks in their necks from talking down to the Community Health Councils and District Health Authorities, and up to Regions and the Department. [Goatly 1997, p.162; from *The Daily Telegraph*]

There is no need to consider or discover parallels for necks, cricks, physical suffering or head-turning: rather, the point of the passage is emotional suffering.

Guidance by Known Generic Mappings

The ATT-Meta project has extensively developed the idea that many of the mappings that are used in metaphor are of general types that do not depend on any specific metaphorical view (conceptual metaphor) such as NOTICING AS SEEING (Barnden and Lee 2001, Barnden 2006, Wallington and Barnden 2006). We say that these mappings arise from mapping principles called *view-neutral mapping adjuncts* (VNMAAs). They are default principles, so their effects can be overridden by more specific information. Amongst these mappings we have included, on the basis of study of many other examples, mappings of (in)ability, difficulty, uncertainty, mental states (including intention to do something),

and modal elements such as possibility and counterfactual-ity. Thus, much of the meaning of the nylon-string example is obtained not by any specific metaphorical view but via these very general mappings, giving the three parallels at the end of the Nylon String subsection. The only view-specific residue is the three parallels preceding those there.

Our analyses of a wide variety of other metaphor examples (in, e.g., Barnden 2006, Wallington and Barnden 2006) strongly suggest that much of the point of a metaphorical utterance is typically conveyed by VNMAAs, rather than view-specific mappings. The view-specific mappings often serve mainly to provide a scaffolding supporting the operation of VNMAAs.

In the neck-crick example, (4), the most important parts of the message are arguably that the managers experience emotional suffering of some sort as a result of their conversations with the other parties mentioned. Here the suffering and the causation are both carried over by VNMAAs in our analysis, not by view-specific mappings.

In the far-reaches example (3), the relevant view-specific mapping is one that connects PHYSICAL OPERATION on physical objects to CONSCIOUS MENTAL OPERATION on ideas. However, in our analysis the meaning for the utterance is that, *presumably*, Anne has a *very low degree* of *ability* to consciously operate mentally upon the idea in question (of her husband Kyle having an affair). The uncertainty in the qualifier “presumably,” the gradedness in the “very low degree” and the “ability” element of this conclusion are all carried by VNMAAs from the source scenario (from the presumed very low ability of Anne’s conscious self to operate *physically* upon the idea). Notice that without some general principle concerning ability one would need a *view-specific* mapping that maps *ability* to physically operate to *ability* to mentally operate, not (just) a mapping of physical *operation* itself to mental operation. This view-specificity of ability mapping would be needed in indefinitely many other metaphorical views as well; and analogous comments apply to matters such as gradedness, uncertainty, causation, emotion and the other things that we postulate are mapped by VNMAAs.

These other things include: (alongside causation and ability, in one broad VNMA), enablement, disablement, ease, difficulty, preventing and helping; (alongside emotional states, in another broad VNMA) value judgments—of goodness, importance, etc.—and mental states (not only of source-scenario agents but also of the understander him/her/itself); the shape of events and processes (e.g., their intermittency, and whether they have a stopping point); time relationships, durations and rates; logical structure; and proper functioning (in the sense of serving a natural or designed function). (This is not complete list.)

The ATT-Meta Approach and System

We now sketch the ATT-Meta approach in slightly more detail, using the Anne/Kyle far-reaches passage, (3). For convenience we recast the initial part of the example as:

- (AKa) In the far reaches of her mind, Anne believes Kyle is unfaithful.

(AKb) To acknowledge the betrayal would mean Anne would have to take a stand.

We assume an understander for whom the “far reaches” of a mind is not stock phraseology allowing simple lexical lookup of a metaphorical meaning, and for whom the phrase denotes a physical-region concept that is unparalleled (has no mapping to an aspect of the mind, according to mappings familiar to the understander). We claim that a reasonable candidate for the main message of (AKa) is that Anne only has a very low degree of ability to consciously think with the idea K of Kyle being unfaithful.

We will first sketch how our general theory applies to the example, and will then comment on more specific points about how the implemented ATT-Meta system works. Other examples of the operation of theory and/or system can be found in Barnden (2001, 2006), Barnden and Lee (2001) and other papers.

ATT-Meta Handling of the Anne-Kyle Example

The only view-specific mapping link we need is a correspondence between physical operation on an idea by an agent’s conscious self and conscious mental usage of it by the agent. We assume the understander knows this mapping link. However, we also postulate that the understander has the following **Ancillary Assumptions** within his/her/its knowledge of metaphor concerning the mind:

- (AS1) If an agent’s mind is being metaphorically viewed as a physical region then (by default) the agent is metaphorically viewed as having a conscious self that is itself a person.
- (AS2) In the source scenario, that metaphorical person is (by default) located in a central subregion of the whole mind-region.
- (AS3) If an agent’s mental usage of a proposition is metaphorically viewed as being physically located in a particular subregion of the mind-space, then the proposition is itself by default metaphorically viewed as a physical object in that same subregion.

Ancillary assumptions such as these serve to flesh out source scenarios with additional details. To our knowledge they are a novel theoretical construct in the explanation of metaphor.

The main steps of understanding according to the ATT-Meta approach are as follows.

(A) From (AKa)’s literal meaning, the premise that *Anne’s believing is physically in the far-reaches of her mind* is created. It is put in a special reasoning space called the **[metaphorical] pretence cocoon**. ATT-Meta casts metaphor as essentially a type of pretence (or counterfactuality).

(B) From that premise it can be **inferred within the pretence cocoon** that Anne’s mind is a physical region.

(C) Using the Ancillary Assumptions (AS1) and (AS2) above, it is defeasibly assumed (within the pretence cocoon) that Anne’s conscious self exists as a person located in the centre of the mind-space. Also, using (AS3), the idea K that

Kyle is unfaithful is defeasibly assumed to be a physical object that is physically located in the far reaches of Anne’s mind-space.

(D) Using commonsense knowledge about physical objects and space, inference in the pretence cocoon now establishes that Anne’s conscious self has (by default) only a very low degree of ability to physically operate upon idea K. This reasoning rests on commonsense knowledge that the far-reaches of a region are distant from the centre of the region and that distance normally affects ability to operate.

(E) From (C), idea K is a physical object. By the mapping link mentioned above, the physical operation in (D) therefore corresponds in reality to conscious mental usage of K. Note that, in ATT-Meta, mapping links map between the contents of the pretence cocoon and reality. (Actually, pretence cocoons can be embedded within other reasoning spaces, rather than directly within reality space. A mapping link operates between a cocoon and its immediately surrounding space.)

(F) However, what is actually inferred in (D) is a very low *degree of ability* to physically operate upon idea K. Our VN-MAs include ones that carry over abilities and degrees identically. These VN-MAs together with the mapping link mentioned in (D) conjointly allow the defeasible inference, in reality space, that Anne only has a very low degree of ability to consciously use the idea mentally, which is our suggested meaning for (AKa).

Note again the non-parallelism inherent in the fact that *no parallel is assumed or created for the far reaches of Anne’s mind* or for the inferential connections between them and other aspects of the source scenario.

Although we have so far portrayed the above reasoning as if it went forward from the initial premise inserted into the cocoon in (A), actually we suppose that it is driven by a backwards-chaining process from an issue raised by context. In the example we take this issue to be that of *what* degree of ability of conscious mental ability Anne has to use idea K. We claim that this issue does arise from (AKb) above.

Current Limitations of Implemented System

The ATT-Meta system works as outlined above on the Anne-Kyle sentence (AK1), with one crucial limitation at present: the system does not yet have an implementation of most VN-MAs. Only the uncertainty VNMA is fully implemented at the time of writing, through the way in which certainty information is propagated through reasoning rules, as described below, notably rules that couch metaphorical mapping links. Thus, although we wish ultimately to use the above mapping link as it stands—i.e., linking physical operation on ideas to mental processing of them, leaving the matters of *degree of ability* to operate/process to be handled by VN-MAs—currently we are forced to use instead a mapping link that explicitly relates degrees of ability to physical operate to degrees of ability to mentally process.

The current implementation provides only the reasoning needed, not a handling of surface forms of utterances. For a specific utterance, a mini-KB is manually created containing hand-constructed logic-based expressions couching the

literal meaning of the utterance, and logical expressions for any additional relevant information that could plausibly be available from context. The system is started by giving it a top-level reasoning query that could plausibly have been derived from context.

More Detail on the ATT-Meta System

The system is a specialized production-rule interpreter implemented in Prolog. Specific knowledge about the world and specific metaphorical mappings and ancillary assumptions are implemented in production rules in hand-created knowledge bases.

Basic Representations Production rules manipulate nodes holding “hypotheses.” Hypotheses are logical forms expressed in an episode-denoting logic broadly similar in spirit to the logical scheme of Hobbs (1990), but for brevity here we will use instead a more conventional logical format. Thus, although John loving Mary is actually represented in the system by the term

```
the-episode(loving, T, john, mary)
```

for some time-interval T, we would instead show the atomic formula `loves(john, mary, T)`, and will omit time arguments.

In essence, hypotheses are things that are at some stage of investigation (perhaps completed). Hypothesis nodes form a complex, dynamic network, in which the main linking consists of support relationships discovered during reasoning (cf. links in justification-based truth-maintenance systems). Each hypothesis node is thought of as lying in one “space,” which in this paper is either the reality space or the pretence cocoon.

Degrees and Uncertainty We use a crude qualitative scale of degrees: `very-low`, `low`, `medium`, `high`, `very-high`, `absolute`. Hypotheses can contain degree specifications of various sorts, one example being:

```
to-at-least-degree(low) [loves(john, mary, T)]
```

Hypothesis nodes contain certainty levels (or, more properly speaking, evidence levels). We use qualitative measures of certainty, for simplicity and naturalness. The certainty level in a node for a hypothesis H is one of: `certain`, `presumed`, `suggested`, `possible`, `certainly-not`. The level `presumed` means that H is a default (a proposition tentatively taken to be true), `suggested` means that there is evidence for H but not (yet) enough to enable H to be a default; `possible` means that the negation of H is not certain; and `certainly-not` means that the negation of H is certain.

Rules and Backwards Chaining The following partially specifies the abstract syntax of rules:

```
<rule> ::= <condition-side> →
          (<rule-qualifier>) <result-side>
<rule-qualifier> ::= suggested | presumed | certain
<result-side> ::= <hypothesis>
<condition-side> ::= <empty> |
                    <condition> {^ <condition>}*
<condition> ::= <ordinary-cond> | <special-cond>
```

```
<ordinary-cond> ::= <hypothesis> |
                  <space-indicator> (<ordinary-cond>)
```

Rules can contain variables as usual. A condition can be the distinguished hypothesis `given`, or `given` surrounded by a space-indicator. Intuitively, a rule containing such a condition represents a given fact. A space-indicator localizes a fact as originating in reality or a pretence.

Rules are applied to queries in a depth-first backchaining mode, starting with the user-supplied top-level query in reality space. When a rule is applied to a query, the sub-queries generated by the conditions are standardly in the same reasoning space as is the query. However, a `<special-cond>` (not detailed above) can cause the resulting sub-query to arise either in an encompassing space (e.g., reality space if reasoning is currently within the pretence space) or an inner space (e.g., dropping down from reality into the pretence space). This is used in rules that effect transfers between spaces. That is, these are the rules that implement metaphorical mapping links.

Certainty Calculus and Conflict Resolution The certainty a rule contributes to its result is the minimum of its own certainty qualifier (`rule-qualifier` in syntax above) and the certainty levels of the hypotheses picked up by the condition part. When multiple rule applications support a given hypothesis node, the maximum of the certainty values they deliver is used. When there is evidence to (merely) level `presumed` for both a hypothesis node and the node for the hypothesis’s negation, a complex conflict-resolution mechanism (see Barnden 2001 for an early version) tries to adjudicate the relative evidence strength in the support network. If one of the nodes wins, it is kept at `presumed` and the other downgraded to `suggested`. Otherwise both are downgraded.

One important use of the conflict-resolution mechanisms is to ensure that unusual circumstances holding in the pretence cocoon can hold sway over contrary knowledge about reality; for example, ensuring that Anne’s mind is held to be a physical space even though in reality it is not.

ATT-Meta, Non-Parallelism and Creativity

We have argued that metaphor often involves non-parallelism but that unparalleled elements nevertheless often play a crucial role in implying the metaphorical meaning. These tenets explain relatively easily how creative (but not entirely novel) metaphor can be understood, thus avoiding substantial analogy-construction overhead. It also illuminates the richness of and scope for creativity in metaphorical expressions, especially because of the possibility that there may in fact be no coherent or useful parallel in the target scenario for all of the source scenario. The ATT-Meta account also explains in detail how rich, subtle and yet important meaning effects can arise from creative, unparalleled source aspects.

The emphasis on non-parallelism does, however, lead to reduced expectations about the extent to which creative metaphor implies a profound new look at the target, since

the basis of the approach is the use of existing mappings (view-specific and/or view-neutral).

Of course, creative metaphor can, under the right circumstances, lead to a profound new look at the target. Although the “far reaches” in example (3) are claimed to have no target parallel, a target parallel is clearly much more desirable, or possibly obligatory, in a sentence such as “The drug shut down the far reaches of her mind,” where the far reaches are part of the main focus rather than being only within a qualification. Then, depending on the understander’s knowledge and views about the actual nature of the mind, he/she/it would be pressured to come up with a correspondent for the far reaches.

ATT-Meta actually allows more creativity and complexity in what is said about the target than may be apparent from the above description. This is because we have included in the approach and system a handling of the compounding (mixing) of metaphor (Lee and Barnden 2001). ATT-Meta freely allows, in mapping from within a pretence cocoon to the surrounding space (usually reality), the co-use of mapping links that happen to be relevant, where these links may intuitively be parts of different (even conflicting) metaphorical views. Another type of compounding supported is serial chaining, explicated in ATT-Meta by nesting of pretence cocoons within each other.

Finally, we can broaden ATT-Meta’s fundamental principles beyond creative metaphorical language to say something about other forms of creativity, and help provide the foundation for computational approaches to them. For one thing, the principles provide a framework for explicating how, in a work of art (visual, musical, or whatever) that is somewhat representational in depicting something outside itself, there may be elements that do not themselves connect directly to aspects of the depicted scenario but merely imply or link to non-explicit elements that do so connect. More generally, the principles explicate how, in any use of analogy in creative processes, it may be that the source side of the analogy contains elements that are essential to the effect of the analogy but that do not themselves take part in it.

Acknowledgments

This work was supported in by grants GR/M64208 and EP/C538943/1 from the EPSRC in the U.K. I have benefited from collaboration with Rodrigo Agerri, Sheila Glasbey, Mark Lee and Alan Wallington.

References

Barnden, J.A. 2001. Uncertainty and Conflict Handling in the ATT-Meta Context-Based System for Metaphorical Reasoning. In *Proceedings of the Third International Conference on Modeling and Using Context*, 15–29. Lecture Notes in Artificial Intelligence, Vol. 2116. Berlin: Springer.

Barnden, J.A. 2006. Consequences for Language Learning of an AI Approach to Metaphor. In J. Salazar, M. Amengual and M. Juan (Eds), *Usos Sociales del Lenguaje y Aspectos Psicolingüísticos: Perspectivas Aplicadas*, 15–57. Palma, Mallorca: Universitat de les Illes Balears.

Barnden, J.A., Glasbey, S.R., Lee, M.G. and Wallington, A.M. 2003. Domain-Transcending Mappings in a System for Metaphorical Reasoning. In *Conference Companion to the Tenth Conference of the European Chapter of the Association for Computational Linguistics*, 57–61. Association for Computational Linguistics.

Barnden, J.A., Glasbey, S.R., Lee, M.G. and Wallington, A.M. 2004. Varieties and Directions of Inter-Domain Influence in Metaphor. *Metaphor and Symbol* 19: 1–30.

Barnden, J.A. and 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 Univ. of Birmingham, UK.

Falkenhainer, B., Forbus, K.D. and Gentner, D. 1989. The Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence* 41: 1–63.

Fass, D. 1997. *Processing Metaphor and Metonymy*. Greenwich, Conn.: Ablex.

Goatly, A. 1997. *The Language of Metaphors*. New York: Routledge.

Glucksberg, S. and Keysar, B. 1990. Understanding Metaphorical Comparisons: Beyond Similarity. *Psychological Review* 97: 3–18.

Hobbs, J.R. 1990. *Literature and Cognition*. Stanford University, Calif.: CSLI Press.

Holyoak, K.J., Novick, L.R. and Melz, E.R. 1994. Component Processes in Analogical Transfer: Mapping, Pattern Completion, and Adaptation. In K.J. Holyoak and J.A. Barnden (Eds), *Advances in Connectionist and Neural Computation Theory, Vol. 2*, 113–180. Norwood, N.J.: Ablex.

Kövecses, Z. 2002. *Metaphor: A Practical Introduction*. Oxford, UK: Oxford Univ. Press.

Lakoff, G. 1993. The Contemporary Theory of Metaphor. In A. Ortony (Ed.), *Metaphor and Thought*, 2nd ed., 202–251. New York: Cambridge Univ. Press.

Lakoff, G. and Johnson, M. 2003. *Metaphors We Live By*, 2nd ed. Chicago: Univ. of Chicago Press.

Langlotz, A. 2006. *Idiom Creativity: A Cognitive-Linguistic Model of Idiom-Representation and Idiom-Variation in English*. Philadelphia: Benjamins.

Lee, M.G. and Barnden, J.A. 2001. Reasoning about Mixed Metaphors with an Implemented AI System. *Metaphor and Symbol* 16: 29–42.

Narayanan, S. 1999. Moving Right Along: A Computational Model of Metaphoric Reasoning about Events. In *Proceedings of the National Conference on Artificial Intelligence*, 121–128. Menlo Park, Calif.: AAAI Press.

Sperber, D. and Wilson, D. 1995. *Relevance: Communication and Cognition*, 2nd ed. Oxford, UK: Blackwell.

Stern, J. 2000. *Metaphor in Context*. Cambridge, Mass.: MIT Press.

Wallington, A.M. and Barnden, J.A. 2006. Similarity as a Basis for Metaphor: Invariant Transfer and the role of VNMA. Technical Report, CSRP-06-02, School of Computer Science, The Univ. of Birmingham, UK.

Wilks, Y. 1978. Making Preferences More Active. *Artificial Intelligence* 11: 197–223.