Metaphor, Semantic Preferences and Context-Sensitivity

John A. Barnden
School of Computer Science; University of Birmingham
Birmingham, B15 2TT, United Kingdom

J.A.Barnden@cs.bham.ac.uk
Tel: (+44)(0)121-414-3816 Fax: (+44)(0)121-414-4281

1 Introduction

In this chapter, the main reference point in Yorick Wilks's work is Wilks (1978). This extends his preference-based semantics (Wilks, 1975) to handling metaphor. Wilks (1978) covers a number of issues that deserve fresh comment and that are central to problems about metaphor that are still unresolved—theoretically, let alone computationally. Also, the legacy of the 1978 work stretched through to, and beyond, the work of Fass (1997) on an approach to metaphor and metonymy based on "Collative Semantics." This was a development from Wilks's Preference Semantics. Fass's system *meta5* implemented an approach to metaphor (and metonymy) derived from Wilks (1978). Wilks himself also continued work on metaphor in conjunction with various researchers, including myself (see, e.g., Ballim, Wilks & Barnden, 1990, 1991; Fass & Wilks, 1983; Wilks, Barnden, & Wang, 1991, 1996).

The Wilks (1978) approach to metaphor is very roughly as follows. Various types of words—notably verbs and prepositions—have "preferences" for the semantic types of words that they semantically liaise with in an utterance. E.g., a verb has preferences for the semantic types of its case-role fillers, such as the agent. These preferences are like the selection restrictions previously suggested by other authors (Katz & Fodor, 1963), but a key difference is that violations of preferences are tolerated, indeed expected, and are used to cause further processing to come up with metaphorical utterance interpretations (as one main possibility). The main example used is the famous "My car drinks gasoline." The initial interpretation of the sentence is a complex internal data structure that can be grossly summarized as *[my-car drink gasoline]*. The verb *drink* prefers an animate agent, but this is violated by *my-car*. Therefore, an encyclopaedic knowledge structure for *car* is searched in an attempt to find something about cars that matches *[my-car drink gasoline]*. This could lead, say, to the item *[a-car-engine use liquid]* being found. As a result, the final interpretation of the sentence is *[my-car use gasoline]*. That is, because of the match found, *drink* has been replaced by *use*.

Wilks (1978) himself said that a person reading the sentence would surmise that the car uses a *great deal* of gasoline, not just *uses* it. However, he said that this is an "idiomatic" part of our interpretation and is not accounted for by his process—indeed, he said that no reasoned basis could be provided for it. We will take up this matter below.

Our concern will not be with the details of the process or with what it is able to achieve on the basis of discovered preference violations. Rather, our concern is to cast the approach as a special case of metaphor interpretation approaches that are largely or wholly *utterance-based*. Fully utterance-based approaches operate by coming up with a metaphorical interpretation based just on the utterance in isolation, without taking context into account. The present paper reacts against this, in the sense of arguing that in many cases the process of metaphor interpretation should in large measure be *contextual-issue*-based, not just utterance-based. That is, interpretation should as far as possible be guided by specific issues raised by the surrounding discourse (or by context in a broader sense). We will see by way of the "ATT-Meta" approach and implemented computer program for metaphor processing (Barnden, 1998, 2001a,b; Barnden *et al.*, 2002, 2003, 2004; Lee & Barnden, 2001a,b) how specifically this can happen, once the specific issues have been extracted from context.

Many metaphor researchers have affirmed the importance of context in metaphor interpretation (e.g., Cameron, 1999; Gibbs & Tendahl, 2006; Giora (1997); Hobbs, 1990; Leezenberg, 1995; Peleg, Giora & Fein, 2001; Stern, 2000). Recently, Carston & Wilson (2005) have suggested the need for a context-based interpretation process roughly on the lines of ATT-Meta's, as a way of dealing with metaphor within Relevance Theory (Sperber & Wilson, 1995). Indeed, Wilks (1978) himself talked about context, pointing out that context could affect how the interpretation process proceeds, thus making it less than purely utterance-based.

However, context-aware work on metaphor has been short on proposals for detailed mechanisms whereby context can affect interpretation. The main exception is probably Hobbs (1990), and indeed the ATT-Meta account of metaphor interpretation is strongly related to his. (For differences, see Barnden, to appear.) In particular, the contextual guidance that the ATT-Meta approach proposes is, as stated above, based on the idea that context raises very specific issues, rather than just providing, say, information about the general subject matter of the discourse. Examples will be given below.

Naturally, interpretation normally has to be based partly on the utterance itself, otherwise it hardly counts as interpretation. It is possible to imagine situations in which an utterance is totally incomprehensible in its own right (whether through, say, corruption or unfamiliar vocabulary) but nevertheless its context is so definite that the meaning of the utterance can be guessed anyway. This fringe possibility aside, interpretation rests in some way both on information from the utterance and information from context.

Although advocating contextual-issue-based approaches, this chapter still affirms the potential for a strong heuristic role for Wilksian semantic preference violations. Violations can still be an important heuristic both in guiding interpretation and in suggesting that the utterance is metaphorical in the first place, even though there are many cases of metaphor that do not violate preferences.

The plan of the rest of the chapter is as follows.

- First there are some observations about the limitations of using preference-violation as a guide to the presence of metaphor (drawing upon other authors but also adding observations of my own). Nevertheless, the discussion preserves the point that metaphorical expressions often do violate preferences.
- We then look at the main general problem with purely utterance-based metaphor processing—the indeterminacy of interpretation—and go on to see how making it partly based on specific, contextually-raised issues can be beneficial.

- In discussing the use of contextual issues, we see in outline the way the author's ATT-Meta approach works. The current implementation of this approach takes a particularly strong line on the usage of contextual issues, namely by using them to generate backwards reasoning that finally connects up with the utterance itself. However, we discuss the possible desirability of having a mix of this backwards, contextual-issue-*driven* style of processing and an utterance-*driven* style that goes forwards from the utterance to finally meet up with the contextual issues.
- We see how the ATT-Meta approach provides an alternative approach to the car-drinking-gasoline example, but also see how preference-violation could still be a useful guide. In addition, we see that the use-*a-great-deal* interpretation mentioned above could in principle be obtained.
- The chapter culminates by sketching an overall framework for metaphor interpretation stretching from the case of stock phraseology, through cases of minor variation of stock phraseology, then through the more open-ended type of metaphor ATT-Meta is mainly addressed at, and on finally to completely novel metaphor. We see that preference violations could be one useful heuristic guide, mainly in the last two cases.

A notable simplification in this chapter, for the sake of brevity, is that we do not attend to a particularly important complication. This is that preference violations can signal phenomena other than metaphor—particularly metonymy, as studied in Fass (1997).

2 Limitations of Preference Violations

Wilks (1978) gives the following expressions as examples of preference-violating metaphorical utterances. The presumed preference-bearing items are italicized and some preference-violating elements underlined. For example, in (a) the verb "taken" bears a preference for its patient being a movable physical object, whereas "line" does not fall in this category.

(a) the line *taken* by the Shadow Cabinet

(b) a Scottish Assembly should be *given* no executive powers

(c) *lead to* the break-up

(d) *break-up of* the United Kingdom

(e) Britain tries to *escape* Common Market

(f) my car *drinks* gasoline

(g) my car *drinks* mud

(h) my car *chews* gasoline

(i) John's new car *runs on* diesel

(j) [John's car] *does* 100 m.p.h.

(k) John *grasped* the idea

(l) I *see* what you mean

(m) An ambulance driver ... *went through* red traffic lights.

For the sake of argument let us agree with Wilks that these examples do plausibly break preferences. The trouble is that variant expressions of many of them, at least, can be constructed that are still metaphorical but for which it is much more difficult to argue that they contain preference-breaking. In the following variants of some of (a–m), the underlining indicates items that correspond to items in (a–m) but that now fail to break the preferences of the italicized verbs.

(a′) the line *drawn* by the Shadow Cabinet

([members of] the Cabinet could literally draw a line, but that is not what's meant here[1])

(b′) a Scottish Assembly should be *given* no explosive device

([members of] the Assembly could literally be given an explosive device, but that is not necessarily what the exhortation is about; rather, "explosive device" could be used metaphorically to refer to something abstract such as a special, powerful law)

(e′) Britain tries to *escape* the prison warders of the Common Market

(using "prison warders" to refer metaphorically to functionaries of the Common Market, but the escaping is not physical escaping as one might do with real prison warders)

(f′) my car *floated down* the street

(a car could literally float on water, and a street could be flooded; but the sentence could also mean that the car motored smoothly and effortlessly down the dry street)

(l′) I *see* the picture you're painting

(could be literal or metaphorical).

The central point here—that a sentence that is clearly metaphorical in context need not show internal signs of being metaphorical, let alone any internal sign as simple as a preference violation—has been pointed out in many ways by numerous authors. However, the point is made especially vividly by doing simple manipulations on utterances that do break preferences or show other internal signs of metaphoricity.

A further example is as follows. It arose in "e-drama", i.e. dramatic improvisation or role-play conducted over computer terminals. Zhang, Barnden, Hendley & Wallington (2006) have studied e-drama in a research project that is centrally concerned with affect-laden metaphor. In an improvised dramatic session concerning school bullying, a character named Mayid has already insulted another character named Lisa by calling her a "pizza," developing a previous "pizza-face" insult. Mayid then says: "I'll knock your topping off, Lisa" – a theoretically intriguing spontaneous creative extension of the "pizza" metaphor (itself possibly a metonymic extension of the earlier "pizza-face" metaphor). The noun "topping" does not violate any reasonable preference of "knock off."

[1] We do not dwell on the complication that "the Shadow Cabinet" may break an agent preference for "draw" and should be interpreted metonymically to refer to its members. This point applies also to example (b′).

In this example, we can also consider what would happen if the understander could attach preferences to a word such as "your." In the bodily-part-of sense of "your," the word "topping" would violate a preference for a body-part. The trouble is there is no clue internal to the sentence that "your" should have this particular sense as opposed to a general possession sense, so that "your topping" could in principle amount to, for example, "the pizza topping you are eating." So there would still be no preference-violation overall.

On the stance adopted in the ATT-Meta project (see Barnden *et al.*, 2004, and other references above), metaphor always involves viewing something as something else that it is not, e.g. viewing a car as an animal, a mind as a physical space, a person as a pizza, Iraq as Vietnam, or whatever. So, there is always a violation of reality in some sense. However, the difference between the semantic types (e.g, inanimate object *versus* animate being) of the two sides of the metaphorical view can not only be narrower than the semantic type-separations envisaged in proposals for such things as semantic preferences, but in fact they can be *arbitrarily* narrow. For example, take "Tuesday is honorary Monday this week" [heard in conversation], uttered because the real Monday was a holiday. Clearly, it would be difficult to strike a semantic type difference between Monday and Tuesday. Other examples of metaphor that involve narrow semantic-type differences are "Iraq is today's Vietnam" "Purple is this year's black" and "Jules Verne is France's H.G. Wells." (Closeness of semantic type is not peculiar to examples of the syntactic form "A is B's C." A variant of the Iraq/Vietnam example could be "We're in Vietnam again" spoken by someone caught up in the Iraq situation.)

Having said all this, metaphorical utterances often do violate semantic preferences that can plausibly be postulated. For instance, in most conceptual metaphors (metaphorical views) of Lakovian style (Lakoff, 1993) a target-domain item and its corresponding source-domain item are of widely different semantic types, making it reasonably likely that the difference will be manifested in a sentence based on the conceptual metaphor. For example, consider the conceptual metaphor of LOVE AS JOURNEY (Lakoff, 1993), where the source domain is physical journeys and the target domain is love relationships. A crossroads in the source domain could then correspond to a difficult decision point in the love relationship. This could come up in a sentence such as "Jack and Jill were at a crossroads in their marriage." A semantic preference of "in" for the two things it relates to be of suitably related types (e.g., if X is in Y and Y is a physical object then Y should be a physical region or container) could then signal the presence of metaphor. Of course, in a context that is clearly about a love relationship the sentence could also have been simply "Jack and Jill were at a crossroads," with no preference violation.[2]

3 Indeterminacy in Purely Utterance-Based Interpretation

Most work on metaphor that attempts to show how, algorithmically, particular information can be extracted from metaphorical utterances takes a largely or wholly utterance-based approach, with little scope for context to guide the process. Within AI this is true not only of Wilks (1978) but also the related system of Fass (1997) and the system of Martin (1990), though not of Hobbs's (1990) approach.

Approaches such as those of Wilks (1978) and Fass (1997) are utterance-based in the especially strong sense of seeking also to *detect signs of metaphoricity* early in the interpretation process and on the basis of the utterance itself. In their approaches, an attempt at metaphorical interpretation is only made if signs of metaphoricity, in the particular form of semantic preference violations, are detected. But it should

[2] In this paper, our term "metaphorical view" and the widely-used term "conceptual metaphor" can be taken to be synonymous, though there are theoretical differences that need not detain us.

be noted that an utterance-based approach need not have this early-detection characteristic: Martin's system proceeds by investigating the literal interpretation and a range of possible metaphorical interpretations on a par with each other, without prior detection of signs of metaphoricity, finally making a choice on the developed representations by means of a scoring mechanism. Thus, the system can only be said to decide on metaphoricity late in the process, and is able to deal, to an interesting extent at least, with the issue of metaphorical utterances that also have a semantically plausible literal interpretation (and/or that have several competing metaphorical interpretations).

But any largely or fully utterance-based approach faces the problem that a given utterance can have a large set of possible metaphorical interpretations. This has been a concern of many metaphor researchers (e.g., Stern, 2000). Naturally, this is the more true the more unconventional the metaphorical phraseology is. And it is perhaps especially a problem for systems that do not rely on previously known source-to-target mappings (such as, for example, those in Lakoff's conceptual metaphors), but rather, as in the systems of Wilks and Fass, find a source/target analogy completely from scratch by comparing structures.

But even Martin's system, which is based on having a stock of known source-target mappings (e.g., they could include the mappings needed for the LOVE AS JOURNEY conceptual metaphor), faces the multiple-interpretation problem even when the utterance does not involve any novel analogy between two domains. This is for several reasons:

1. The nature of the target domain may not be evident from the utterance itself, as in "Jack and Jill were at a crossroads" above. There may be several possible mappings to several different target domains from the source subject matter contained in the utterance. For example, the physical world is used as source domain for metaphorically talking about many different things – relationships, time, money, mental states, etc. An attempt could therefore be made to use context to select a target domain, but several possibilities might be supported by context.

2. Much metaphor is concerned with conveying evaluations, emotions and other affective attitudes about the target subject matter, rather than making cold propositional points about it (see, e.g., Mio, 1997; Musolff, 2004; Vervaeke & Kennedy, 2004). This happens by the transfer to the target of evaluations, emotions, etc. that are somehow related to source-domain items in the sentence. There are easy cases of this, where the "somehow related" is a very direct relationship and there is no real competition between different possible affective qualities. This would arise in, perhaps, "My son's room is a cess-pit" where negative feelings about cess-pits are highly salient. However, it is trickier to infer the relevant affective quality in sentences like "Lisa is a pizza" (example from the previous section) and "Veronica is being chased by publishers." In the Lisa example, the appropriate affective connotation in context is negative, even though pizzas have strong positive affective qualities for many people. In the publishers example, the appropriate affective connotation in context could well be positive even though *physical* chasing, which is presumably the source of the metaphor, might by default be taken to have a negative quality. Even with "cess-pit" the matter is not completely clear-cut, as cess-pits would be a major technological improvement for people who did not have sanitation at all. Such is the affective indeterminacy of metaphor that different contexts could make *opposing*, not just *different*, types of affect inferrable from one and the same metaphorical utterance in different contexts.

3. The proliferation of possible interpretations is exacerbated by many linguistic factors, but most relevantly here by the fact that metaphorical views (conceptual metaphors) can be mixed together in the same sentence (Lakoff and Turner, 1989; Lee & Barnden, 2001; Wilks, Barnden & Wang, 1991),

and that metonymy and metaphor can be mixed together (Fass, 1997). An example of the mixing of metaphorical views is "This idea crystallized the nebulous mental meanderings that had plagued me ..."[3] where mental processes are viewed as animate, gas-like and disease/germ-like. An example of the mixing of metaphor and metonymy is "The whole of my childhood rushed through my head like an electric train"[4] where (arguably) it is not the childhood as such that metaphorically rushed through her mind, but rather the memories of her childhood, so that there is also a metonymic step from the childhood to the memories. This example could in principle be re-analysed as pure metaphor, with the mind viewed as a physical region that can contain not just ideas but also the entities those ideas are about. But the availability of this possibility just underscores the point about interpretations proliferating.

As we have already noted, many authors have suggested that contextual information is important in metaphor interpretation or have demonstrated by psychological experiment that context affects interpretation, and the underlying assumption is that context can guide metaphor interpretation in the appropriate direction. However, the mechanistic details of how context helps in processing have hardly been explored. We turn to these matters in the next two sections.

4 A Way in Which Context Can Help

The ATT-Meta approach supposes that the context of any utterance, including a metaphorical one, often creates *queries* in the mind of the understander, or, what comes to the same thing, puts certain specific *issues* in focus. The understander uses such queries or issues to control the interpretation of the utterance at hand. In this view, utterances are placed in discourse for specific purposes related to how they connect to other utterances, and the understander benefits by trying to to divine what those purposes are.

The matter is best conveyed by example.[5] Suppose discourse contains the sentence

(1) "John is a tank"

and the understander does not know any metaphorical sense of the word "tank" that could apply to people, and therefore has to consider another sense of the word to guide the process of metaphorical understanding. For the sake of brevity, let us assume that that other sense is the military-tank sense. Even with this restriction, the sentence could be getting at a variety of different things, such as: John is square and heavy; John destroys things; John is tough; John tramples over things — where, moreover, the destruction, toughness or trampling could be physical or abstract.

A purely utterance-based approach to understanding (1) would blindly have to take qualities of military tanks, such as size, heaviness, inexorableness, ability to withstand attack, and powerfulness, and transfer them in some form to apply potentially to John. The hope would be to find that one or more of these transferred qualities could plausibly apply to him.

[3] From Sheila Dyan, *Love Bites,* London: Hodder & Stoughton, 1992, p.48.

[4] Heard on the BBC Radio 4 programme *Desert Island Discs* on 20 June 2003.

[5] The material in the rest of this section is based on a section of Barnden *et al.* (2004).

But it is hardly likely that (1) would be uttered in a context that gave no independent clue as to what the sentence was trying to get at. More plausibly, (1) would appear in a specific context. A possibility is:

(1′) "Most of my colleagues get dispirited when they're criticized, but John's a tank."

The first clause here raises the issue of the ability to tolerate criticism. The word "but" suggests a contrast between the two clauses. The understander can therefore pose an internal query to him/her/itself such as:

(2a) **Is John able to tolerate criticism well?**

For definiteness, let's assume that the understander knows an ARGUMENT AS WAR metaphorical view (Lakoff & Johnson, 1980), and has ready access to a mapping link between *military attack-withstanding* and *criticism-tolerating*. So, on encountering (2a) one thing that the understander can do is to use this mapping to translate (2a) into military terms:

(2b) **Is John able to withstand military attack well?**

The source-domain query (2b) can then be answered in the affirmative within the source domain using the datum that John is a tank.[6]

The information from the sentence that John is a tank, which is a military object, can conceivably be used proactively to boost the relevance of the above mapping link in preference to other mapping links, if any, that do address criticism-tolerating but that do not involve the military domain. In other work (Barnden, 2006) we cast doubt on the utility of domains in defining the very notions of metaphor and metonymy, and we therefore depart from the concentration on domains in approaches such as Lakoff's. However, this does not preclude the possibility that pieces of information are tagged as being about particular rough domains such as warfare and that this information is used heuristically to facilitate metaphorical processing.

For our argument to carry through, it is not necessary for the context to raise the issue of toleration of criticism in quite such an explicit or precise fashion as the first clause of (1′) does. Rather, we make the general assumption that context provides relatively determinate information about what issue the utterer is addressing in the metaphorical utterance at hand, at least in those cases where that utterance itself is indeterminate about that.

But of course it could happen that context does not raise the issue at all or only very implicitly raises one, or raises many issues, so that it was not clear which query or queries the understander should address. However, the situation even in these cases is no worse for the ATT-Meta account here than it is for metaphor theories which do not have any account at all of how context could help with metaphor understanding. Furthermore, when context fails to raise clear issues, a human understander would presumably be unsure about what the metaphorical utterance such as (1) was conveying.

And there is no need to require context to suggest a definite *answer* to any query raised. Although context *may* suggest that one answer is more to be expected than another, the ATT-Meta approach does not assume this happens. The interpretation process should in any case generally seek for evidence both for and against any particular hypothesis that arises. (This is because most information and reasoning that crops up in discourse interpretation is uncertain.) Therefore, even if context were to provide some evidence for a particular hypothesis, this could be overridden by contrary evidence coming from the sentence.

[6] Use of the word "well" in (2a) is a symptom of a general phenomenon whereby metaphor is often about matters of degree rather than yes/no issues. We will see more of this in the discussion of example (3) below.

Also, we do not need to assume that the issue-raising context arises *before* an utterance such as (1)—it could occur afterwards instead. It would be reasonable for an understander to adopt a strategy whereby if preceding context does not raise any discernible, specific issue the understander postpones full interpretation of the sentence until after examining succeeding context.

We have used invented tank examples in order to isolate the issues of interest, but appropriate examples can readily be found in real discourse. Real-discourse examples using the noun "rock" are analysed in detail in Barnden *et al.* (2004). One of these examples is:

> "Okay. My husband has always been very involved with the children, although he works a lot of hours. He spent more time than he usually does with them. Obviously, I wasn't around, or I was sick, but *he was a rock.*" [italics added]

Notice that in this example, just as in our tank example, the context raises a specific issue partly through the conjunction "but" (in "but he was a rock"). This appears to be a common phenomenon in metaphorical discourse. Some of Hobbs's (1990) examples also rely on a "but." And "but" occurs also in the following example from a magazine article,[7] part of which we will proceed to analyze in detail:

(3) *In the far reaches of her mind,* Anne knew Kyle 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.

The analysis of this example in Barnden & Lee (2001) and Barnden *et al.* (2004) is as follows. The discourse chunk involves the metaphorical views of MIND AS PHYSICAL SPACE and of IDEAS AS PHYSICAL OBJECTS. Assume that the meaning of the segment "In the far reaches of her mind, Anne knew Kyle was having an affair[.]" to be in part that Anne only had a very low degree of conscious awareness of the idea that Kyle was having an affair. The succeeding segment "to acknowledge the betrayal," together with the "but," can readily give rise to a target-domain query such as

(4a)**To what degree was Anne able to operate in a conscious mental way on the idea of Kyle having an affair?**

Let us assume that IDEAS AS PHYSICAL OBJECTS includes a correspondence between the physical operation on ideas, by the agent's conscious self, with conscious mental operation on them by the agent. Thus (4a) can be transformed to create the query

(4b)**To what degree was Anne's conscious self able to operate physically on the idea of Kyle having an affair?**

This query then controls the understanding, within the source domain, of the significance of the qualifier

(4c)"In the far reaches of her mind"

[7] In Linden Gross, "Facing up to the Dreadful Dangers of Denial," *Cosmopolitan, 216*(3), USA ed., March 1994. Italics in original.

in (3). This qualifier indicates indirectly *a very low level of ability by Anne's conscious self to operate physically on the idea.* This answer to (4b) is transferred, by the above-mentioned correspondence, to the target domain to become the conclusion that Anne only had a very low degree of ability to operate in a conscious mental way on the idea. This conclusion is the answer to query (4a).

Qualifier (4c) indicates a very low level of ability by Anne's conscious self to operate physically on the idea because, in applications of MIND AS PHYSICAL SPACE, the conscious self of the person is implicitly viewed as being a person located in a main part of the physical space, presumably distant from "far reaches." This distance in turn implies a very low degree of ability of the conscious self to physically interact with the idea.

This interpretation process is of the sort advocated in the ATT-Meta theoretical approach. In fact, given the original query (4a), the process is fully implemented in the ATT-Meta system. An overall picture of the major reasoning steps taken is given in Figure 1 (queries themselves are not shown, except for the top query, 4a). The process is explained in considerable detail in Barnden & Lee (2001). The "pretence space" in Figure 1 is a special computational environment in which inferential consequences of the ostensible meaning of the utterance (e.g., the proposition that Anne's state of knowing really does have a physical location within her mind, which really is a physical space) can be teased out without risk of contaminating or being contaminated by reality. For purposes of the present chapter, the inference within the pretence space can be taken to be within the terms of the source subject-matter, and the reasoning within the reality space to be within the terms of the target subject-matter. However, there is in fact no restriction on what subject matter arises in each space.

An important point to notice from the Figure, for the purposes of the next section, is that the reasoning can be split into three broad aspects:

1. the reasoning within the pretence space (i.e., that joins up the utterance's ostensible meaning with the source side of one or more metaphorical mapping links);
2. the reasoning consisting the actual use of one or more mapping links;
3. and the reasoning in the reality space that connects the target sides of the utilized mapping links with the specific context-derived queries.

This list implies no specific temporal ordering, however, for reasons that we do not go into here.

5 Context-Drivenness in ATT-Meta

The ATT-Meta system embodies a particularly strong way of using contextual issues that arise. The issues take the form of reasoning queries like the TOP QUERY in Figure 1 (expressed of course in a formal internal representation language), and these are used to generate a backwards-chaining reasoning process that eventually meets up with the ostensible meaning of the utterance in the pretence space. That is, given a query such as (4a), ATT-Meta looks for facts or IF-THEN rules in its knowledge based that could provide a value for the variables in the query (the degree variable in the case of 4a), or, if there are no variables, supports or refutes the hypothesis. If it finds such rules, then their IF parts are used to create sub-queries,

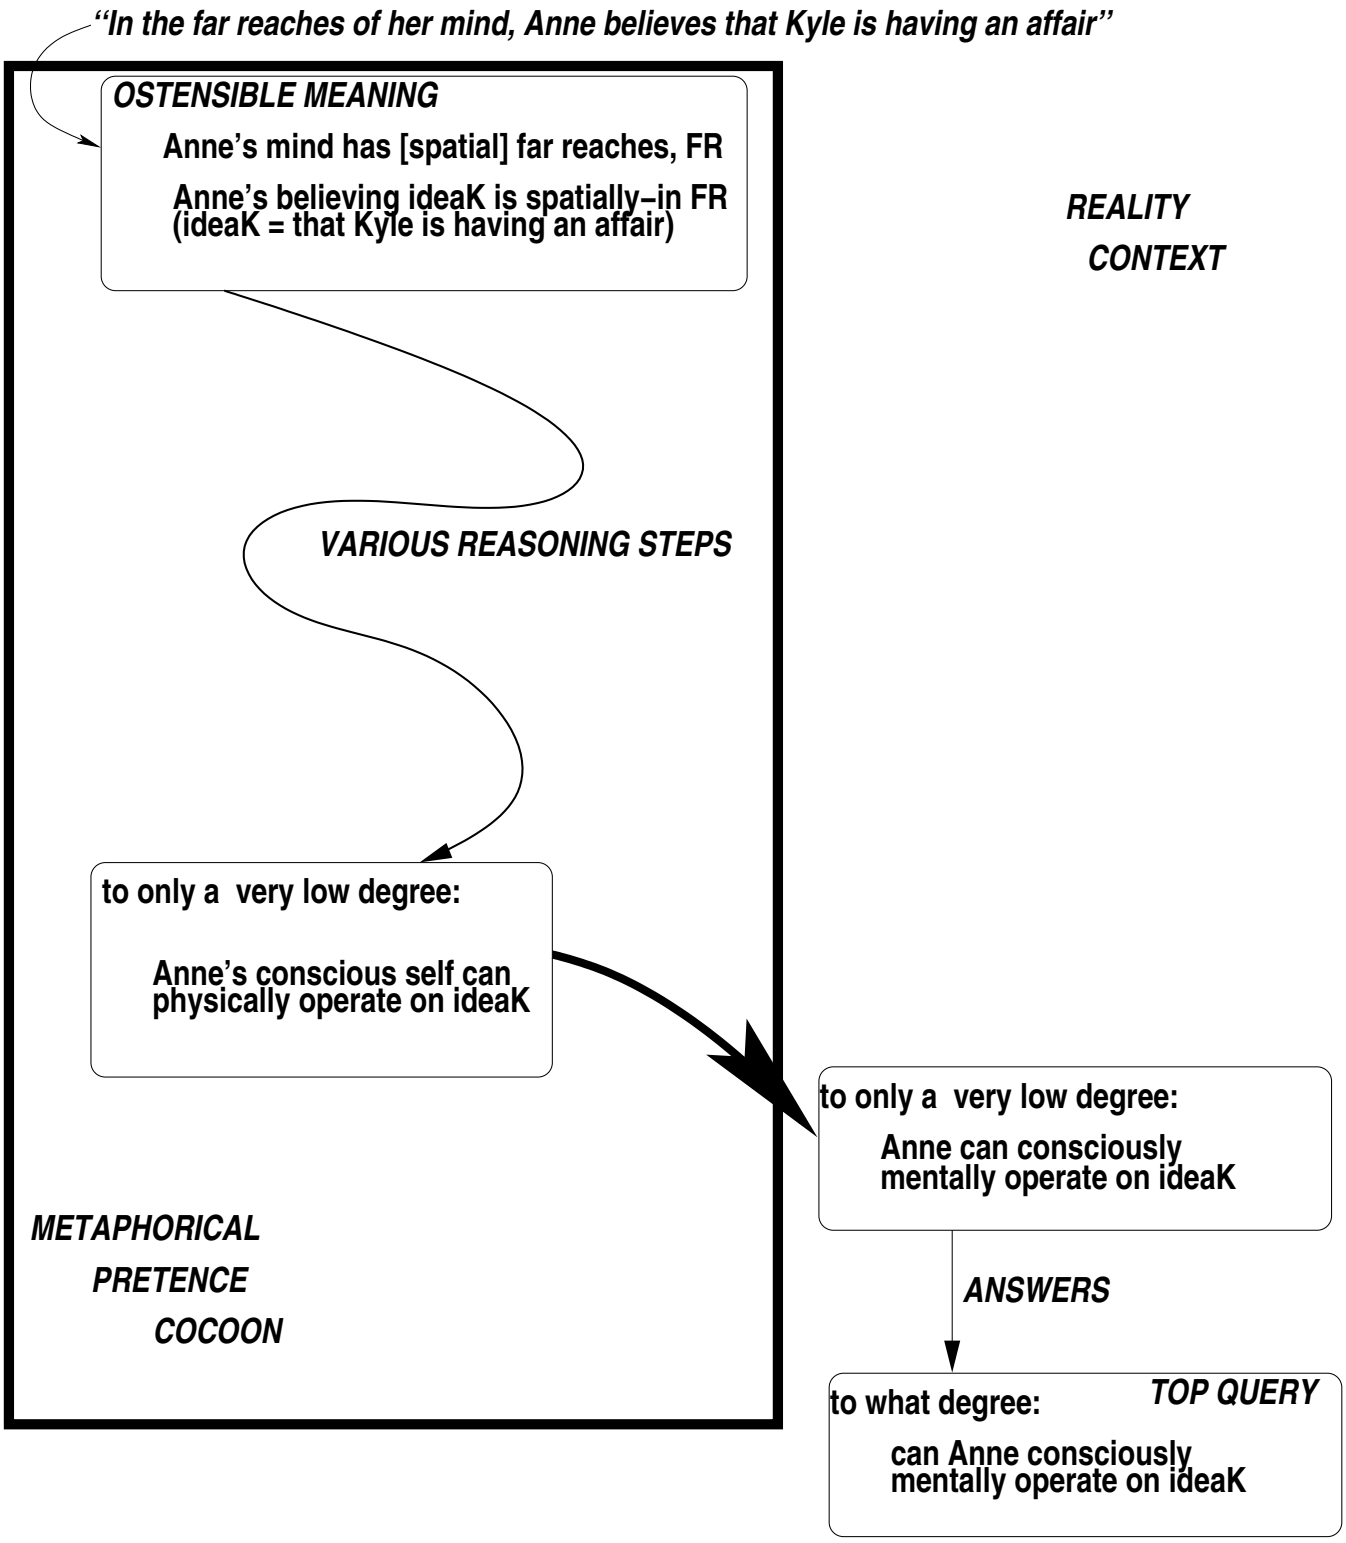

Figure 1: **Overall shape of processing of an initial part of the Anne/Kyle example, (3) in text. The thick arrow shows the action of mapping. The statements within the diagram are English glosses of expressions in ATT-Meta's internal representation scheme.**

and so forth. Since metaphorical mapping links are themselves cast as IF-THEN rules, the creation of, say, query (4b) from (4a) fits naturally and flexibly into the backwards chaining process. We can call the overall process context-*driven* rather than merely context-*based* because of the backwards, query-directed nature of the reasoning.

However, the ATT-Meta approach has always recognized that some measure of *forward* chaining from the utterance's ostensible meaning may be needed or desirable (see, e.g., Barnden & Lee, 2001). Such processing could be called utterance-*driven* rather than just utterance-based. What mix of utterance-driveness and contextual-issue-drivenness is desirable is partly a practical matter of how to connect the ostensible meaning to the contextual issues in the most effective and efficient way. Only much further empirical work on testing the approach on specific examples, with realistically sized knowledge bases, will reveal the best mix or suggest how the mix should be adjusted to suit different circumstances. The system and theoretical approach have been tested on many examples (see Barnden, 2001b,c; Barnden & Lee, 2001; Lee & Barnden, 2001) but the knowledge bases used have not been large—for example, just under 60 rules or so were used in the experiments on the Anne/Kyle example ("*In the far reaches of her mind,* Anne knew Kyle was having an affair" from (3)).

Thus, it is worth emphasizing that there are two possible theses that should be distinguished:

1. Metaphor interpretation should, when possible, not just be utterance-based but also contextual-issue-based.

2. When metaphor interpretation has a contextual-issue-based aspect, this aspect should be handled through contextual-issue-driven processing rather than utterance-driven processing.

Much suggestive evidence can be provided that context of metaphor does often raise specific issues and that these issues can effectively be used to guide interpretation, thereby supporting thesis 1. There is less evidence that thesis 2 is true even though the current version of the ATT-Meta *system* has adopted it as a working hypothesis. The overall ATT-Meta theoretical approach, on the other hand, is not committed to relying wholly on contextual-issue-driven processing.

There is nevertheless a general argument that suggests that contextual-issue-drivenness is often, and perhaps normally, the method of choice. The argument appeals to the conjecture that metaphor generally casts a relatively abstract and ill- and/or sparsely-understood subject matter (the target) in terms of a relatively concrete and well/richly-understood subject matter (the source) (see, e.g., Lakoff, 1993). We will call this the conjecture of *source superiority on concrete richness.* When this superiority exists, there will typically be more choices to make for any given reasoning step about the source subject matter than is the case for the target subject matter. The remainder of the argument is as follows, appealing to the three reasoning types listed at the end of the last section.

Because of the extra richness in the source subject-matter, it is more difficult to go forward in an utterance-driven way from the ostensible meaning of a metaphorical utterance towards the source side of mapping links (reasoning of type 1) than it is to go backward from the contextual queries to the target sides of mapping links (reasoning of type 3). In addition, in the contextual-issue-driven case, once a query has been transformed by going backwards over a mapping link to land within the source subject-matter, the sub-queries that thereby arise are from only a *subset* of all known mapping links, and furthermore these queries arise from *specific applications* of the links. By contrast, utterance-driven reasoning that is aimed

at connecting up with source sides of mapping links does not know which mapping links will turn out to be useful, and has no specific applications of those mapping links to work with.

Having said this, it is still the case that, because source superiority on concrete richness is at best only a general tendency, there may be occasions on which it is practical to go forwards from ostensible meanings to mapping links without (yet) having any guidance from context. This is especially so to the extent that reasoning within sources may often be constrained by quite specific stereotypical "scenarios" of the type that Musolff (2004) has discussed (cf. scripts, etc. in AI). For instance, he argues that in political metaphors casting the EU as a marriage, there are standard scenarios of getting married, being unfaithful, or getting divorced etc. that tend to be used, rather than the full panoply of possible knowledge about marriage in general being deployed.

A further consideration is as follows. As we noted above, one purpose of many metaphors is to convey affect, e.g. evaluations of or emotions about the target situation described. It appears that such affective elements are often carried over identically in metaphor, irrespective of the particular metaphor-specific mapping links involved (such as a mapping link between lovers and travellers). Thus, strong disgust towards cess-pits carries over in "My son's room is a cess-pit." Now, given that affect often does carry over in this way and may indeed be part or all of the point of the metaphorical utterance, it makes some sense for an interpretation process to try to reason forwards from the utterance to find affective consequences, irrespective of the specific context at hand.

We saw above, though, that conflicting affective consequences may arise from a given utterance, and used this as one argument about the *dangers* of purely utterance-based processing. Thus, there is a balancing act needed between the dangers of proliferating interpretations by not using contextual guidance in inferring affective qualities and the benefits of having affect as a useful source of guidance.

Affect is just one of sizable set of properties that the ATT-Meta project takes to be carried over identically in metaphor. A provisional set of properties the ATT-Meta project has been working with is listed in Barnden & Lee (2001) and Barnden *et al.* (2003), and includes, for example, temporal structure, causation relationships, dis/enablement relationships, and purposes. Since the mapping principles involved are neutral with respect to any particular metaphorical view, they are dubbed as being *view-neutral mapping adjuncts* (VNMAs). What we will call here the "Affect VNMA" is the principle that, first of all, if the understander judges a source-domain item S to have affective quality Q, and there is a target-domain item T corresponding to S, then T also has quality Q, by default. (All effects of VNMAs are defeasible.) A second aspect of the Affect VNMA is that if an agent A in the source-domain scenario judges source-domain item S to have affective quality Q, and both A and S have corresponding target-domain items B and T, then (by default) B judges T to have quality Q.

VNMAs are loosely related to a hierarchy of invariant characteristics that Carbonell (1982) claimed tended to be carried over in metaphor, though VNMAs are not regimented in a hierarchy and are different in detail.

The non-affect VNMAs, such as one that carries over causation relationships from source to target, could provide an end-point to within-pretence reasoning, much as the Affect VNMA could. Indeed, analysis of some examples of metaphor (Barnden, 2001b) suggests that in many cases *most* of the useful information coming from a metaphor is carried by VNMAs as opposed to view-specific mapping links. View-specific links can often be largely or wholly confined to providing a scaffold that allows the relevant applications of VNMAs. This phenomenon occurs in the next section.

6 Gas Guzzler Tamed

In order further to illustrate the ATT-Meta approach, we address the central example in Wilks (1978), namely

(5) My car drinks gasoline.

It is also a prominent example in Fass (1997), and we use the treatment there, in the meta5 system, as a contrast to the way the ATT-Meta approach would proceed on the example. The meta5 system can interpret the sentence as meaning "My car uses gasoline" essentially by finding an analogical match, from scratch, between the following two knowledge items: animals drink liquids; cars use gasoline. The meta5 system has no prior knowledge of any particular metaphorical views. However, we can arguably reanalyze the example more naturally within the ATT-Meta framework.

Plausibly, ordinary English users possess a metaphorical view (conceptual metaphor) of MACHINES AS CREATURES (as indeed Fass, 1997:p.318, points out). Utterances such as "my radio is dead," "my car has life in it still," "he killed the engine," "a middle-aged toaster," and "an aggressive lawn-mower" are mundane and easily understandable. We assume that as part of the view, a machine's running corresponds to a creature's biological activity. *This is the only view-specific mapping we need in order to be able to deal with (5).*

The proposed reasoning process is sketched in Figure 2. It has not been implemented in the ATT-Meta system, because it relies heavily on some VNMAs which have not yet been implemented. From the ostensible meaning of the utterance and source-domain general knowledge, it can be (defeasibly) inferred in source-domain terms that gasoline helps the car to be alive (biologically). But, by default, being alive enables the creature to be biologically active. Therefore, gasoline helps the car to be biologically active. Since being-biologically-active maps to machine-running as mentioned above, and as helping is mapped across identically by a VNMA (the one that also deals with causation, enablement, ability, etc.), the pretence-space hypothesis that the gasoline helps the car to be biologically active can be transformed into the reality-space hypothesis that the gasoline helps the car to run. Further reasoning sketched in Figure 2 could now produce the conclusion that the car *uses [up]* gasoline.

The production of these target-domain conclusions does not require any mapping of the drinking itself to be created. Of course, an understander *could* go on to do the extra work of mapping drinking itself to, say, the process of a car having gasoline put in it or of the engine using the gasoline. Furthermore, to say that the drinking maps to *using*, as opposed to something else such as having gasoline *put in,* seems an unwarranted extra commitment. The ATT-Meta approach does not make that commitment since the inference to the *car-uses-gasoline* proposition is defeasible (because the approach is based on uncertain reasoning), and defeat of it would still leave us with the useful proposition that *gasoline helps the car to run.*

As pointed out above, Wilks (1978) himself says that a person reading the sentence would get a connotation that the car uses *a great deal of* gasoline, not just uses it, and says that this connotation is not accounted for by the process be outlines. Fass (1997:p.192) himself makes a similar brief observation. For purposes of applying the ATT-Meta approach, we can frame the desired connotation as being that the car uses gasoline relatively quickly. Compare "the blotting paper drank up the ink." Also the entries in Webster's *Third New International Dictionary* suggest at least moderate rapidity of ingestion. But, because an act of drinking is (normally) moderately fast, on a rate-of-change scale concerning ordinary human activities, a

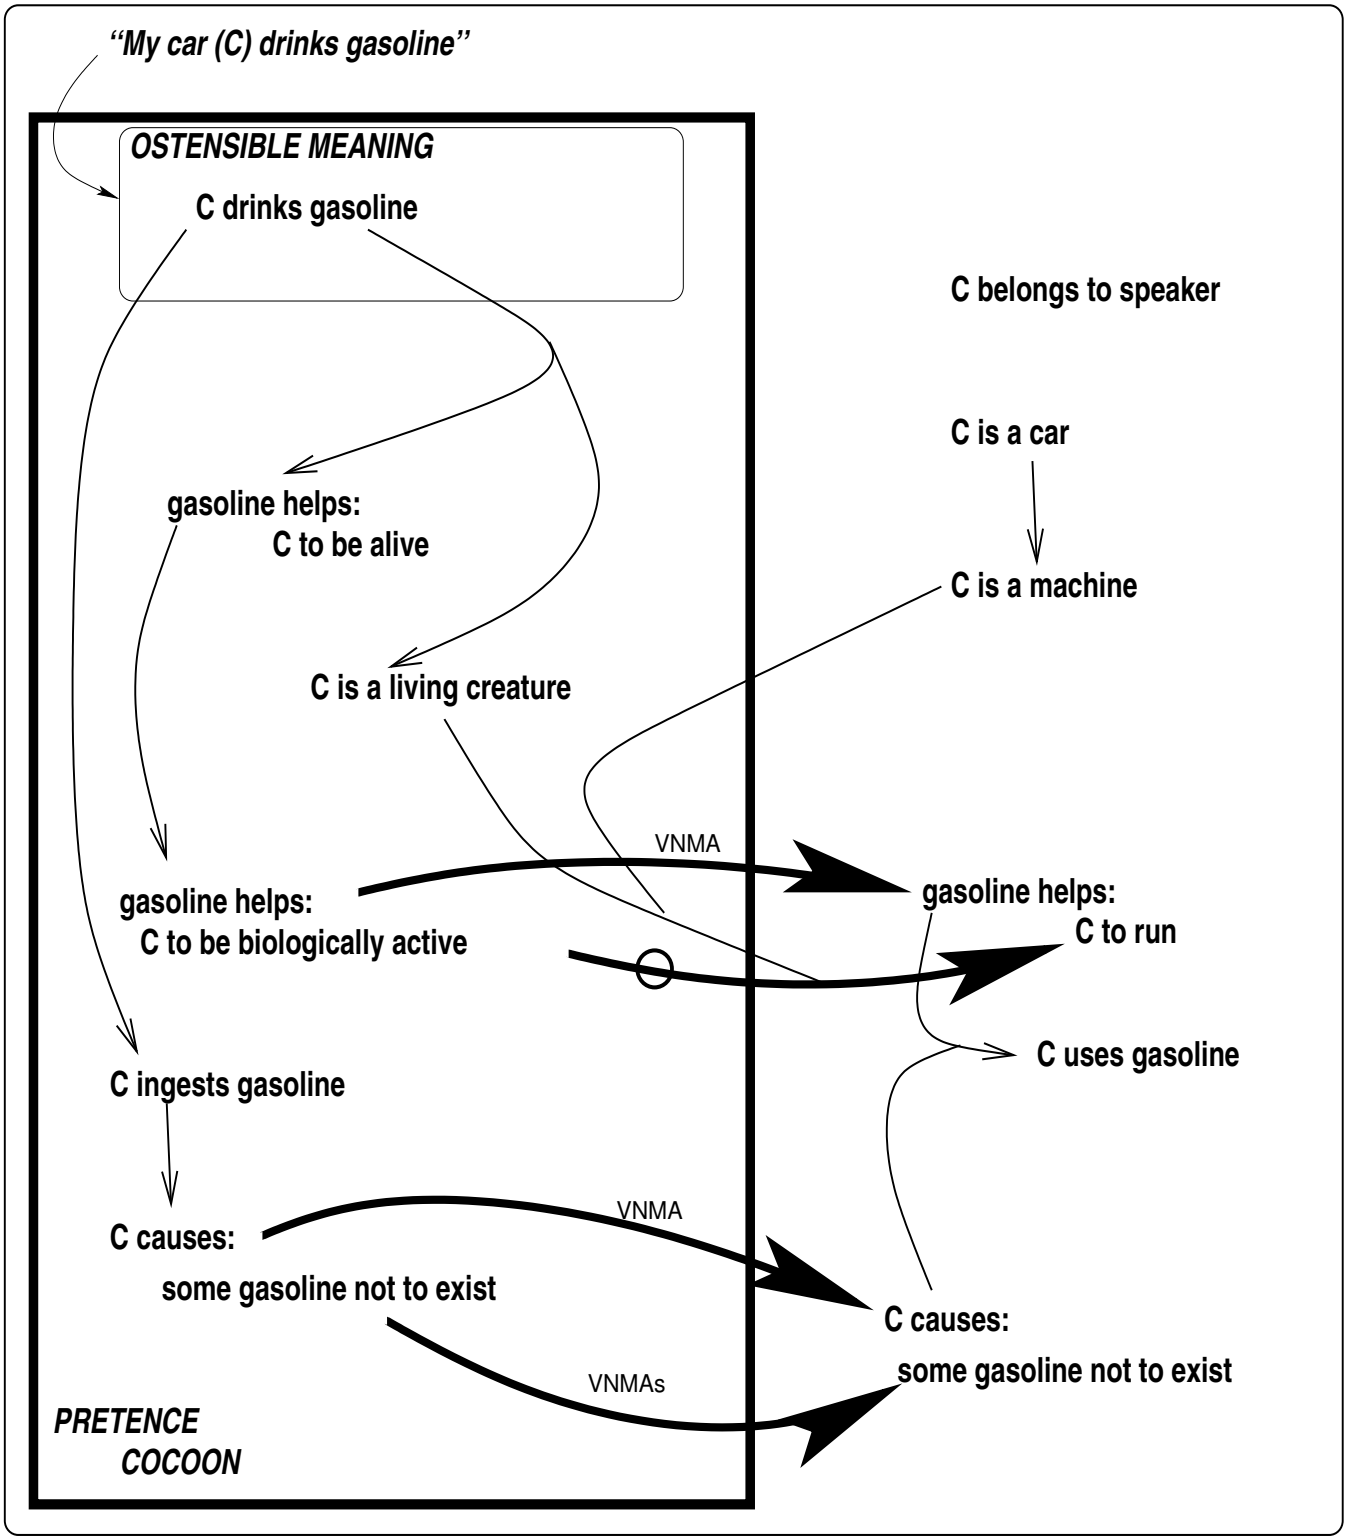

Figure 2: **Showing how the approach could deal with The car-drinking-gasoline example, (5) in text. A thick arrow labelled VNMA or VNMAs shows the action of one or more VNMAs. A thick arrow marked with a circle shows the action of a mapping relationship specific to the particular metaphorical view, MACHINES AS CREATURES. The statements within the diagram are English glosses of expressions in ATT-Meta's internal representation scheme.**

use of a VNMA that deals with temporal rate (Barnden & Lee, 2001; Barnden *et al.*, 2003) would allow the ATT-Meta approach to conclude that the car's use of gasoline is (probably) moderately fast, relative to the normal speed of consumption.

In order for meta5 to be able to come up with the connotation that the car's consumption is relatively quick, it seems that it would need to already have in the target domain a representation of cars using gasoline relatively quickly, because otherwise there would be nothing in the target domain to be found to be analogous to the source-domain situation. But surely the point of the connotation in question is the *exceptional* circumstance that the particular car mentioned in the sentence uses gasoline relatively quickly.

The example, as treated by ATT-Meta, provides an instance of the point made at the end of the previous section about the importance of VNMAs as opposed to view-specific mapping links. In Figure 2, most of the pretence-to-reality transfer (i.e., source to target transfer, in effect) is done by VNMA applications. The only view-specific transfer is that of being-biologically-active to mechanically-running, and no special, extra correspondence apart from this needs to be discovered in order to interpret the sentence.

This section has not yet explained the way contextual-issue-basedness or contextual-issue-drivenness could come in. If someone said to you "My car drinks gasoline" in the absence of any prior context, you could well be forgiven for not being sure what the person was getting at. Potentially, the world-knowledge that cars use gasoline could by itself generate a contextual-issue-driven process conforming to Figure 2. However, it is surely more likely that there would be prior linguistic context, as perhaps in "I had to fill up this morning again. My car drinks gasoline." (And it would be reasonable to expect there to be intonational emphasis on the "drinks," or for the sentence to be varied to "My car really/just drinks gasoline.") The first sentence has the implication that the car has used a lot of gasoline recently. This could be framed as an issue to guide the interpretation of the second, metaphorical, sentence. As "using a lot recently" is tantamount to "using up fast," we have a suggestion here about how the interpretation could be guided.

Another possible scenario, involving a change in the example, is if person A says "My car uses leaded gasoline" and B replies "Oh, mine drinks unleaded [gasoline]." Assuming that a contrast is hinted at here, the issue of whether B's care uses *leaded* gasoline or not could be raised to guide the interpretation of the second sentence. Potentially, this could lead to investigation of the possibility that B's car uses unleaded gasoline, as this would imply that it probably does not use leaded. The issue of speed of consumption could well not be raised.

Finally, although the overall ATT-Meta approach to the car-drinks-gasoline example is very different from the Wilks and Fass accounts, preference violations could still give a clue as to what mappings are involved. Thus, the fact that "drinks" prefers an animate creature as agent suggests that the car is being viewed as a creature, in turn suggesting the involvement of the metaphorical view of MACHINES AS CREATURES.

7 An Overall Picture

By way of a conclusion, the overall picture of how to deal with different types of metaphorical utterances that preceding parts of this chapter suggest is as follows.

(A) *Stock metaphorical phraseology* (completely fixed conventional metaphor). Such phraseology can have

its target-domain meanings listed in a lexicon, as has often been observed in the literature. Included here are not only particular words and phrases (potential examples are "see" in the sense of understand, "probe into" in the sense of abstractly investigate, and "build castles in the air," with of course inflection of the verbs allowed) but also templates that have internal gaps that need to be filled (e.g., "in the recesses of [someone's] mind" and "at the back of [someone's] mind").

Assuming that such stock items are listed in the understander's lexicon, there is no need to detect any metaphoricity, or to detect preference-breaking or other anomalies, although a phrase may in fact contain some anomaly such as violating semantic preferences of an included word or phrase. The recesses/back-of-mind cases are examples of this.

(B) *Minor, open-ended variants of stock metaphorical phraseology* obtained, for example, by (i) replacement of words by synonyms and (ii) adding modifiers such as adjectival/adverbial words or phrases. With reference to the examples in category (A), an illustration of (i) would be the various possibilities in "*construct/erect/elevate/...* castles in the air," and illustrations of (ii) would be "see *dimly*," "probe *deeply* into," "in the *dark, murky recesses* of [someone's] mind" and "at the *very* back of [someone's] mind." Moon (1998) provides an extensive corpus-based study and discussion of such and other forms of variation.

Something to be stressed here the open-ended quality of the variation: any way of, for example, conveying the activity of building (not only by a single verb, but perhaps by a creatively constructed phrase) could replace the verb "build" in the castles example, and any way of conveying visibility problems could be used in place of the "dark, murky" modifier in the recesses example to get a similar effect. Such variation would not tax the understanding powers of a competent speaker of English. Naturally, a particular variant of a stock item might itself happen to be a stock item in its own right, but not all possible variants can be.

Given the open-endedness, which prevents the lexicon-listing approach used for category (A) being applied to the variants in general, it is typically necessary to reason about the effect of the variations in terms of the source domain. That is, special metaphorical processing is needed, perhaps of ATT-Meta style, to some appreciable extent.

The modifiers in variants include ones like "very" that are neutral as to which metaphorical view is involved, whereas modifiers like "dark" are specific to certain source subject-matters. But even the use of "very" in "the idea was at the *very* back of [someone's] mind" requires metaphorical processing that involves the source domain. Certainly, the adjective "very" introduces a general intensifying atmosphere, but it is important to work out exactly what it is that is intensified: namely, the subsidiarity of the mentioned idea. The intensification is not of the other important aspect of the meaning of the phrase, namely that the idea is still *present* to conscious thought in some way. It is difficult to see how to get the correct intensification without reasoning in terms of the source domain.

In order for the understander to realize that the special metaphorical processing for category (B) is needed, it would be advantageous for the understander to detect the metaphoricity when this can readily be done. And indeed, given that the discussed variations of subtype (B)(ii) (modifier addition) are syntactically minor, even when added modifiers are in themselves syntactically complex, metaphoricity detection is likely not to be a major problem for subtype (B)(ii), as modifiers can be stripped out for purposes of comparison to stock items. It is less clear how to deal with B(i), however, especially when a word in the stock item has been replaced by a complex phrase. But, assume that to some useful extent it is possible to determine underlying stock items from encountered variants. Then, if stock items in the lexicon have their metaphorical senses tagged as being metaphorical, the metaphoricity of the variant is a default inference from the metaphoricity

of the underlying stock item. So, processes of stripping out modifiers, looking for synonyms, etc. could uncover metaphoricity for variation types (B)(ii) and some cases of (B)(i). Clearly, though, preference violations would be a useful adjunct for determining metaphoricity, especially when the underlying stock item cannot readily be determined.

In addition, stock items in the lexicon could have their metaphorical senses annotated with information about what mapping links are involved. Then, when it is noticed that an expression in discourse is a variant of a stock item, the mapping links can be accessed from the lexicon by the metaphorical processing, without their needing to be worked out afresh.

Type (A) itself allows some variability by means of holes in templates. However, the variability in (A) is a matter of obligatory choice of fillers for specific holes, with the non-hole parts absolutely fixed aside from morphological inflection, whereas the variability in (B) is optional, relatively unregimented syntactically, and generally applicable to all elements of the base phrase.

(C) *Metaphorical utterances that do not fit in (A) or (B) but nevertheless can be analysed as relying on familiar metaphorical views, i.e. familiar sets of mappings.* A central form of this phenomenon can be called *map-transcendence.* Map-transcendence arises when the source-domain scenario involves an element that is not itself directly mappable by any mapping the understander knows, even though some metaphorical view the understander knows is involved in the utterance. Map-transcendence is similar to the notion of metaphor that exploits "unused" parts of the source domain (Lakoff & Johnson, 1980; and see discussion in Grady, 1997), and is sometimes referred to as "extension" of metaphor. Relatively minor forms of map-transcendence are exhibited by the cases of category (B) that involve, for example, source-domain synonyms of words in the stock item or source-domain-specific modifiers, where the synonyms or modifiers are not mappable by a known mapping.

But category (C) includes also a more thorough-going open-endedness of metaphorical phraseology, which is able to exploit the resources of the source domain more fully. For example, "in the far distant reaches of [someone's] mind" may not be a stock template or a variant of one, and the "reaches" of a mind may not be mappable by any known mapping. However, the utterance rests on a very familiar view of a mind as a physical region. Another example is "Company A gobbled up company B and spat all its managers out," resting on a familiar view of companies as voracious creatures but creatively transcending it by bringing in the spitting-out.

For map-transcending metaphor, reasoning via the source domain is needed, and again the ATT-Meta approach is a candidate. Although metaphoricity may in general only be finally decided late in processing, heuristic hints such as preference violations, if present, could be a useful guide as to whether the metaphorical processing is likely to be profitable.

In addition, preference violations could give a clue as to what mapping links are involved. We saw in the car-drinking example that the fact that "drinks" prefers an animate creature as agent leads to the suggestion that the metaphorical view of MACHINES AS CREATURES is involved.

(D) *Completely novel metaphor.* That is, metaphor that does not rely on known metaphorical mappings (other than metaphor-unspecific mapping principles such as VNMAs—see the section above on context-drivenness in ATT-Meta). Completely novel metaphor is probably quite rare, even in poetry (Lakoff & Turner, 1989) but it needs to be accounted for. We can make the following three observations.

(i) A first conjecture about completely novel metaphors is that a large proportion of them are *image* metaphors, i.e. metaphors that merely consist of a similarity of shape or visual appearance. By extension we could include here other types of perceptual similarity, such as acoustic. Of the examples of metaphor in Goatly (1997) that I judge to be novel, many are, plausibly, image metaphors. An example might be "This pencil is a snake" if the pencil is flexible and wiggly, or "My car is a shrub" if the car is covered with branches and leaves after an off-road driving experience. It would appear that such metaphor needs to be dealt with by special perception-aware similarity processing, and is quite likely not to have any meaning other than the perceptual similarity itself. But one unresolved problem is how to distinguish image metaphor from non-image metaphor, if this is possible at all, in general. For example, a car could be a decrepit one stationed in a garden, and "My car is a shrub" could potentially be used to get at the point that birds build nests in it, without the car having much perceptual similarity to a shrub. Again, it could all be a matter of the issues that arise in the context. In addition, there is no reason why a metaphor should not combine a perceptual-similarity aspect with other aspects.

(ii) Putting aside the special case of image metaphor, the existing computational approaches that are most relevant to completely novel metaphor are ones like those of Wilks (1978), Fass (1997), an aspect of Hobbs's (1990) proposal, and analogy-finding models such as SME (Falkenhainer, Forbus & Gentner, 1989) that discover structural analogies, rather than more holistic perceptual similarities, between source and target knowledge structures from scratch.

However, if it is right that much of the point of a metaphor is often carried by VNMAs, then there is no reason to think that this does not apply to completely novel metaphor in particular. For example, with "my car is a rotten banana," one (defeasible) source-domain inference that might be drawn is that the banana is disgusting. The affect VNMA would then lead to the (defeasible) inference that the car is disgusting in some way. In suitable contexts this could an appropriate, and perhaps even a highly important, inference. If VNMAs are important for novel metaphor, then it follows that establishing complex structural analogies between source and target is concomitantly less important.

(iii) The processing of completely novel metaphor is likely to benefit on many occasions from detection of preference violations, as it is reasonable to think that the metaphor will often bring together widely disparate subject matters.

The above four categories of metaphor, (A) to (D), form an overall, rough framework for seeing how different types of metaphor handling may fit into an overall picture of metaphor. The division into types of metaphor is not itself remarkable, and is not greatly different from categorizations produced by other authors (e.g., Goatly, 1997). The point to note is the particular association of algorithmic methods to the categories. The categorization is language-user-relative in the sense that where a particular utterance lies in the framework is partly dependent on what the particular utterer or understander has in in his/her/its lexicon and what metaphorical mappings he/she/it possesses. For example, whether a particular expression is common enough to count as a "stock" item, and what degree of stock-ness is used to warrant inclusion in a lexicon, are user-relative matters.

The framework accords varying degrees of usefulness to preference-violation detection. This usefulness is potentially for two things: (1) noticing that there is metaphoricity, and hence for giving more weight that there might otherwise be for special metaphorical processing such as the analogy-finding of Wilks or Fass or the the inferential processes of ATT-Meta; and (2) suggesting the specific metaphorical mappings that may be in play. Purpose (1) is potentially important for metaphor categories (C) and (D) and to a lesser extent for (B), but is not important for (A). Purpose (2) is often important when purpose (1) is. It could also

be important for category (B) if the lexicon entries in question did not specify mapping links or state what rough domains or semantic types the mappings are between.

But it should not be forgotten that heuristics other than preference violation are potentially helpful for metaphoricity detection as well. For instance, Goatly (1997) lists a number of morphological, lexical and general phraseological clues to metaphoricity.

Acknowledgments

This research was supported in part by grants GR/M64208 and EP/C538943/1 from the Engineering and Physical Sciences Research Council (EPSRC) of the UK, and grant RES-328-25-0009 from the Economic and Social Research Council and EPSRC. I have benefited from collaboration with the other members of the Figurative Language Research Group at the School of Computer Science, University of Birmingham, U.K.—Sheila Glasbey, Mark Lee, Alan Wallington and Jane Zhang—and with former colleagues Steve Helmreich, Eric Iverson and Gees Stein, as well as Yorick Wilks, while I was at the Computing Research Laboratory, New Mexico State University. During that period the precursor to the metaphor research represented in this article was supported by NSF grant IRI-9101354.

References

Ballim, A., Wilks, Y. & Barnden, J. (1990). Belief ascription, metaphor, and intensional identification. In S.L. Tsohadzidis (Ed.), *Meanings and Prototypes: Studies in Linguistic Categorization.* New York: Routledge, Chapman & Hall. pp. 91–131.

Ballim, A., Wilks, Y. & Barnden, J.A. (1991). Belief ascription, metaphor, and intensional identification. *Cognitive Science, 15* (1), pp.133–171.

Barnden, J.A. (1998). Combining uncertain belief reasoning and uncertain metaphor-based reasoning. In *Procs. Twentieth Annual Meeting of the Cognitive Science Society,* pp.114–119. Mahwah, N.J.: Lawrence Erlbaum Associates.

Barnden, J.A. (2001a). Uncertainty and conflict handling in the ATT-Meta context-based system for metaphorical reasoning. In V. Akman, P. Bouquet, R. Thomason & R.A. Young (Eds), *Procs. Third International Conference on Modeling and Using Context,* pp.15–29. Lecture Notes in Artificial Intelligence, Vol. 2116. Berlin: Springer.

Barnden, J.A. (2001b). 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. (2001c). Application of the ATT-Meta metaphor-understanding approach to various examples in the ATT-Meta project databank. Technical Report CSRP–01–02, School of Computer Science, The University of Birmingham, U.K.

Barnden, J.A. (to appear). Metaphor and artificial intelligence: Why they matter to each other. To appear in R.W. Gibbs, Jr. (Ed.), *Cambridge Handbook of Metaphor and Thought*, Cambridge University Press.

Barnden, J.A. (2006). Metaphor and metonymy: A practical deconstruction. Technical Report CSRP–06–1, School of Computer Science, The University of Birmingham, U.K.

Barnden, J.A., Glasbey, S.R., Lee, M.G. & Wallington, A.M. (2002). Reasoning in metaphor understanding: The ATT-Meta approach and system. In *Procs. 19th International Conference on Computational Linguistics*, pp.1188–1193. San Francisco: Morgan Kaufman.

Barnden, J.A., Glasbey, S.R., Lee, M.G. & Wallington, A.M. (2003). Domain-transcending mappings in a system for metaphorical reasoning. In *Conference Companion to the 10th Conference of the European Chapter of the Association for Computational Linguistics* (EACL 2003), pp.57–61. Association for Computational Linguistics.

Barnden, J.A., Glasbey, S.R., Lee, M.G. & Wallington, A.M. (2004). Varieties and directions of inter-domain influence in metaphor. *Metaphor and Symbol, 19*(1), pp.1–30.

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.

Cameron, L. (1999). Operationalising 'metaphor' for applied linguistic research. In L. Cameron & G. Low (Eds), *Researching and Applying Metaphor,* 1–28. Cambridge, U.K.: Cambridge University Press.

Carbonell, J.G. (1982). Metaphor: an inescapable phenomenon in natural-language comprehension. In W. Lehnert & M. Ringle (eds), *Strategies for Natural Language Processing,* pp.415–434. Hillsdale, N.J.: Lawrence Erlbaum.

Carston, R. & Wilson, D. (2005). Metaphor and relevance: The 'emergent property' issue. Talk delivered at *New Directions in Cognitive Linguistics: First UK Cognitive Linguistics Conference,* University of Sussex, Brighton, U.K., October 2005.

Falkenhainer, B., Forbus, K.D. & Gentner, D. (1989). The Structure-Mapping Engine: algorithm and examples. *Artificial Intelligence, 41* (1), 1–63.

Fass, D. (1997). *Processing metaphor and metonymy.* Greenwich, Connecticut: Ablex.

Fass, D. C. & Wilks, Y. (1983.) Preference semantics, ill-formedness, and metaphor. *J. Association for Computational Linguistics, 9* (3 & 4), 178–187.

Gibbs, R.W., Jr. & Tendahl, M. (2006). Cognitive effort and effects in metaphor comprehension: Relevance theory and psycholinguistics. *Mind and Language, 21*(3), 379–403.

Giora, R. (1997). Understanding figurative and literal language: The graded salience hypothesis. *Cognitive Linguistics, 8*(3), 183–206.

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

Grady, J.E. (1997). THEORIES ARE BUILDINGS revisited. *Cognitive Linguistics, 8*(4), pp.267–290.

Hobbs, J.R. (1990). *Literature and cognition.* Stanford University, CA: CSLI Press.

Katz, J. & Fodor, J. (1963). The structure of a semantic theory. *Language, 39,* pp.170–210.

Lakoff, G. (1993). The contemporary theory of metaphor. In A. Ortony (Ed.), *Metaphor and Thought*, 2nd edition, pp.202–251. New York and Cambridge, U.K.: Cambridge University Press.

Lakoff, G. & Johnson, M. (1980). *Metaphors we live by.* Chicago: University of Chicago Press.

Lakoff, G. & Turner, M. (1989). *More than cool reason: a field guide to poetic metaphor.* Chicago: 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.

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.

Leezenberg, M. (1995). *Contexts of metaphor.* ILLC Dissertation Series, 1995–17, Institute for Language, Logic and Computation, University of Amsterdam, The Netherlands.

Martin, J.H. (1990). *A computational model of metaphor interpretation.* San Diego, CA: Academic Press.

Mio, J.S. (1997). Metaphor and politics. *Metaphor and Symbol, 12*(2), pp.113–133.

Moon, R. (1998). *Fixed idioms and expressions in English.* Oxford, U.K.: Clarendon Press.

Musolff, A. (2004). *Metaphor and political discourse: Analogical reasoning in debates about Europe.* Basingstoke, UK: Palgrave Macmillan.

Peleg, O, Giora, R. & Fein, O. (2001). Salience and context effects: Two are better than one. *Metaphor and Symbol, 16* (3&4), pp.173–192.

Sperber, D. & Wilson, D. (1995). *Relevance: Communication and cognition* (2nd Edition). Oxford: Blackwell.

Stern, J. (2000). *Metaphor in context.* Cambridge, MA and London, UK: Bradford Books, MIT Press.

Vervaeke, J. & Kennedy, J.M. (2004). Conceptual metaphor and abstract thought. *Metaphor and Symbol, 19*(3), pp.213–231.

Wilks, Y. (1975). A preferential, pattern-seeking, semantics for natural language inference. *Artificial Intelligence, 6*, pp.53–74.

Wilks, Y. (1978). Making preferences more active. *Artificial Intelligence, 11*, pp.197–223.

Wilks, Y., Barnden, J. & Wang, J. (1991). Your metaphor or mine: belief ascription and metaphor interpretation. In *Procs. 12th Int. Joint Conf. on Artificial Intelligence* (Sydney, Australia, Aug. 1991). San Mateo: Morgan Kaufmann. pp. 945–950.

Wilks, Y., Barnden, J. & Wang, J. (1996). Your metaphor or mine: belief ascription and metaphor interpretation. In B.H. Partee & P. Sgall (Eds), *Discourse and Meaning: Papers in Honor of Eva Hajičová,*pp.141–161. Amsterdam/Philadelphia: John Benjamins.

Zhang, L., Barnden, J.A., Hendley, R.J. & Wallington, A.M. (2006). Exploitation in affect detection in improvisational e-drama. In *Procs. 6th International Conference on Intelligent Virtual Agents. Lecture Notes in Computer Science, 4133*, pp.68–79. Springer.