METAPHOR AND SYMBOL, *16*(1&2), 29–42

Reasoning About Mixed Metaphors Within an Implemented Artificial Intelligence System

Mark G. Lee and John A. Barnden
School of Computer Science
University of Birmingham

The phenomenon of mixed metaphor has traditionally been viewed as secondary to the understanding of straight metaphors. This article suggests that such an assumption is detrimental to long-term research. It is claimed that the same kinds of reasoning and knowledge structures involved in understanding straight metaphors are also required in understanding mixed metaphors and that mixing is a central phenomenon. Therefore, any theory of metaphor must be able to account for mixing. To this end, the article provides an analysis of both parallel and serial mixed metaphors that has been implemented in an artificial intelligence system that is already capable of reasoning about straight metaphors.

Mixed metaphors are often regarded as humorous or as cases of defective speech. Consider the following example, quoted by Fowler (1908) in his guide to writing style:

1. "*This, as you know, was a burning question; and its unseasonable introduction threw a chill on the spirits of all our party.*"

The question is metaphorically "hot." However, its introduction makes the party's spirits "cold." Despite this contradiction, the sentence can be understood to mean that the question was somehow controversial and its inappropriate introduction saddened the emotions of the party members. Furthermore, it is plausible that most readers would not even consider the disparity of hot questions causing cold reactions because the two pieces of temperature information could be separately

Requests for reprints should be sent to John A. Barnden, School of Computer Science, University of Birmingham, Edgbaston, Birmingham B15 2TT, England. E-mail: J.A.Barnden@cs.bham.ac.uk

mapped to provide connotations relevant in the tenor domain. Owing to these conflicts, such an example makes blending the two metaphors in a single space (as might be advocated by Turner & Fauconnier, 1995) difficult. In this article, we argue that given a mix involving two familiar metaphors for which there are established mappings between the tenor and vehicle domains, the default is that processing is done in two separate "cocoons" (special computational environments).

Under mixed metaphor we include not only examples that might be regarded as obvious cases of conflict, bad style, or humor, but also examples that include graceful combinations of metaphors, such as the following sentence to be examined later: *"One part of John hotly resented the verdict."* This combines a view of John as made up of subagents and a view of agents' emotional states as things that can have temperature. There has been relatively little research done on the topic of mixing due to an assumption that the problem of straight metaphors should be dealt with completely before attempting to tackle the more complex case of mixing. In this article, we argue that this assumption is detrimental to progress because mixed metaphors rely on the same conceptual knowledge as straight metaphors and can, therefore, provide valuable insight into the processes and representations underlying metaphorical reasoning. Moreover, this article makes the following claim: The reasoning processes and data structures involved in understanding mixed metaphors are identical to those used in understanding straight metaphors. Therefore, any current theory of metaphor should (at least in principle) be extensible to deal with mixing.

To this end, we describe some initial work done with ATT-Meta (Barnden, 1997a, 1998; Barnden & Lee, 1999), a computational model of metaphorical comprehension, to handle various types of mixing. We also reprise an earlier claim for the need for extensive within-vehicle reasoning (Barnden, Helmrich, Iverson, & Stein, 1996) and the use of conversion rules to filter the relevant connotations of a particular metaphor.

MIXED METAPHOR DISTINCTIONS

It is possible to distinguish two types of mixed metaphor: parallel mixes and serial mixes. In a *parallel mixed metaphor,* the tenor (A) is seen partly through an *"A as B"* metaphor and partly through another metaphor, *"A as B′."* B and B′ are in general different domains, but may overlap. Also, different aspects of A may be involved in the two metaphors. In a *serial mixed metaphor* (commonly called a *chained metaphor*), the tenor (A) is seen as a vehicle (B), which is in turn seen as a different vehicle (C). For example, consider the following two mixed metaphors:

2. *"John's research wounded the theory's shaky foundations."*
3. *"One part of John hotly resented the verdict."*

The utterance in the Example 2 manifests two familiar conceptual metaphors: *"ARGUMENT AS WAR"* (Lakoff & Johnson, 1980), if we construe war in a broad way, and *"THEORIES AS BUILDINGS"* (Grady, 1997; Lakoff & Johnson, 1980). The former is manifested in the verb *wound;* literally, physical living beings wound physical living objects, so both the research and the basic assumptions of the theory are being viewed as physical living beings. However, the theory is also being viewed as a building, and its basic assumptions as physical foundations of it. Following the definition given earlier for parallel mixes, the following domains are involved:

A: Domain of theories, ideas, arguments, and so on.
B: Domain of living beings.
B′: Domain of buildings.

Shaky foundations in a building suggest that the building itself might collapse; therefore, if a theory is a building, then its shaky foundations may cause the entire theory to collapse or, literally, be refuted. The sentence is best unraveled by treating the different metaphors (*"A as B," "A as B′"*) separately, because there is conflict between the theory's basic assumptions being viewed as living beings and being viewed as foundations of a building.

The utterance in Example 3 also manifests two familiar conceptual metaphors: *"MIND PARTS AS PERSONS"* (Barnden, 1997b; see also Lakoff, 1996, on metaphors of self and Lakoff, 1993, on *"IDEAS ARE ENTITIES"*) and *"ANGER IS HEAT"* (e.g., Lakoff & Johnson, 1980). In the *"MIND PARTS AS PERSONS"* metaphor, the mind is composed of different person-like parts that may have different beliefs, emotions, and personalities. Mentioning that one part of John resented the verdict suggests that there exists more than one part and that some other part of John did not resent the verdict. Moreover, the part of John referred to resented the verdict *"hotly."* In the *"ANGER AS HEAT"* metaphor, anger is seen as heat. Therefore, the part of John that resented the verdict did so with anger. Hence, if we assume that under *"MIND PARTS AS PERSONS"* the emotions of parts are emotions the whole agent tends to have, John tended to be angry. Following the definitions just given, the following domains are involved:

A: Domain of John's mental and emotional states and processes.
B: Domain of people and natural language communication.
C: Domain of heat.

Example 3 is a serial mixed metaphor. The *"ANGER AS HEAT"* metaphor (B as C) acts on the *"MIND PARTS AS PERSONS"* metaphor (A as B) to directly affect its metaphorical meaning. It is not possible to isolate the two metaphors as in Example 2. It is a mind part that is viewed metaphorically as behaving or feeling

hotly, and the mind part is in turn an aspect of a metaphorical view of the topic of John's mind, rather than being directly an aspect of that topic.

Notice also that there is a subtle distinction we wish to capture: Either one part of John is resenting the verdict and one part is not, and the part resenting is doing so *"hotly,"* or both parts of John are resenting the verdict but only one is doing so *"hotly."* Our intuitions suggest that the former interpretation is the default and we only provide a detailed analysis for this interpretation. However, our treatment is sensitive to such distinctions (as is our computational implementation) and is capable of reasoning about such uncertainties.

We gave for Example 2 a parallel mixing interpretation, but we could conceivably give it a serial mixing interpretation instead: According to this, the theory's foundations are viewed as a building as before, and the foundations of the building are then viewed as an animate being, perhaps because of a more generally applicable *"INANIMATE OBJECT AS ANIMATE BEING"* metaphor. However, we view serial mixing as more complex than parallel, so that unless there are pressing reasons to the contrary we prefer to adopt a parallel analysis in case of ambiguity. In any case, given that the parallel analysis is at least a possible one, it is useful to have an account of it.

In contrast, we claim that it is difficult to postulate a parallel reading of Example 3. This is because such a reading would need to view some real component or aspect of John's mind to be viewed as both a subperson and hotly resenting the verdict. However, we claim that the use of *"MIND PARTS AS PERSONS"* carries no implication that the subpersons are mapped to real components or aspects of the mind. Rather, properties of the parts, individually or in conjunction, are mapped to properties of the whole mind. Thus, *"one part of John's mind"* has reference only in vehicle domain B, not the tenor domain.

A COMPUTATIONAL ACCOUNT

The examples discussed in this article are implemented within the ATT-Meta model of metaphor comprehension. We only detail here the concepts relevant to the current work, but further details can be found in works by Barnden (1997a, 1998) and Barnden and Lee (1999).

ATT-Meta is an artificial intelligence (AI) system capable of both simulative reasoning about beliefs and metaphorical reasoning. Reasoning is done by the use of back-chaining rules of inference that allow differing degrees of certainty. Nested reasoning spaces are allowed to facilitate simulation of other agents and metaphorical reasoning. Two types of nested space are maintained: simulation-pretense cocoons and metaphor-pretense cocoons. Simulation-pretense cocoons are used to model the beliefs of other agents. Metaphor-pretense cocoons are a special type of simulation-pretense cocoon in which the agent modeled is a hypothetical agent who is assumed to believe the manifested metaphor is literally true. For the remainder of this article we are concerned only with metaphor-pretense cocoons.

Knowledge of different domains is encoded in sets of facts and rules that apply to a particular domain. Because conceptual metaphors involve a mapping from one domain (the vehicle) to another (the tenor), ATT-Meta uses conversion rules that explicitly map propositions from one domain to another. ATT-Meta has a small set of conversion rules for each metaphorical view it knows about, and it holds knowledge about the vehicle domain of each such view.

Therefore, any conventional metaphor can be defined by constructing a set of rules to represent the vehicle domain plus a suitable conversion rule or a small set of such. Understanding proceeds by creating a metaphor-pretense cocoon (reasoning space) where the manifested metaphor is taken as literally true, then mapping implications to the tenor domain via conversion rules. Figure 1 shows the cocoons involved in parallel and serial mixing schematically.

ATT-Meta is distinctive in that it licenses extensive within-vehicle reasoning in addition to more common, within-tenor reasoning and vehicle-to-tenor mapping. Rather than simply mapping a correspondence from the vehicle to the tenor and then performing inference to fully understand the connotation of an utterance, ATT-Meta favors extensive inference in the metaphor-pretense cocoon prior to mapping in an effort to produce information that can be mapped by conversion rules. This gives any conversion rule the important function of filtering out nonrelevant parts of a particular metaphor. This is essential for metaphor-pretense spaces to be chained in a sensible manner when dealing with difficult examples such as in Example 3.

Parallel Mixed Metaphors

As discussed earlier, Example 2 relies on two familiar conceptual metaphors. Considering the former, *"ARGUMENT AS WAR,"* we assume that ATT-Meta is familiar with the metaphor and so knows the following correspondence:

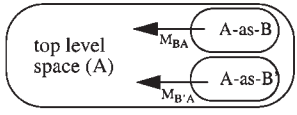

A: Domain of theories (tenor).
B: Domain of living beings (vehicle).
B': Domain of buildings (vehicle').
M_{BA}: Mapping from B to A.
$M_{B'A}$: Mapping from B' to A.

A. Schematic representation of a parallel mix

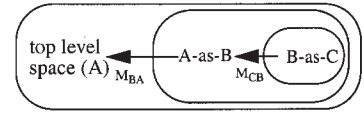

A: Domain of John's mental states (tenor).
B: Domain of people/communication (vehicle).
C: Domain of heat (vehicle).
M_{BA}: Mapping from B to A.
M_{CB}: Mapping from C to B.

B. Schematic representation of a serial mix

FIGURE 1 Arrangement of cocoons (reasoning spaces) for parallel and serial mixed metaphors. The outermost space is the system's top-level reasoning space, in which the tenor domain (A) is reasoned about nonmetaphorically. Other spaces are nested within this. For instance, the cocoon marked as A-as-B is the one for reasoning under the pretense that A really is B.

Metaphorical Correspondence: (Fight-Argue)
Physically damaging an argument/idea/theory/and so on that is being viewed as a battle participant corresponds to establishing faults in the argument/idea/theory/and so on.

This correspondence (and similarly, others later) is couched as a set of conversion rules in ATT-Meta. In addition, suppose ATT-Meta believes the following commonsense rule concerning *"living beings"*:

(Wounding): If X wounds Y then X physically damages Y.

Second, we assume that ATT-Meta is familiar with the metaphor *"THEORIES AS BUILDINGS"* and, as part of this familiarity, knows the following correspondences:

Metaphorical Correspondence: (Instability)
If X is a theory that is being seen as a building then X being unstable corresponds to X being implausible.

Metaphorical Correspondence: (Foundations)
If X is a theory that is being seen as a building then Y being the foundations of X corresponds to Y being the basic assumptions of X.

In addition, ATT-Meta has the following commonsense rule about real buildings:

(Stability): If X is a building and its foundations are shaky, then X is unstable.

Given these mappings and rules, it is possible to infer the connotations that John's research established faults in the theory's basic assumptions and that the theory was implausible by the steps of inference shown in Figure 2. Notice the preceding analysis allows both instances of metaphor to be reasoned about separately. As we see in the next section, serial mixes are more complex.

Serial Mixed Metaphors

As discussed earlier, Example 3 relies on two familiar conceptual metaphors. Considering the former, *"MIND PARTS AS PERSONS,"* we assume that ATT-Meta is familiar with the metaphor and knows the following correspondence:

Metaphorical Correspondence: (State-Tendency)
If person P is viewed as having a part X that is a person, then if X has mental state S then P has a tendency to have state S.

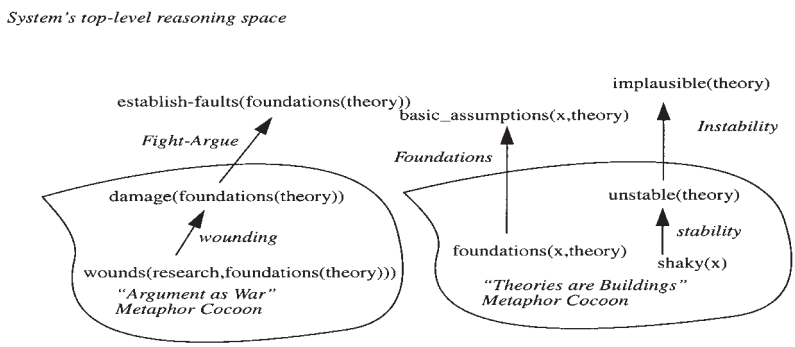

FIGURE 2 Dealing with the parallel mixed metaphor *"John's research wounded the theory's shaky foundations."* Labels next to arrows are rule names used in the text. Notice there is reasoning within the two metaphorical cocoons as well as mapping actions across cocoon boundaries. That reasoning is termed *within-vehicle reasoning* in the text. It can be much more complex and extensive than that shown.

Regarding the second metaphor, *"ANGER AS HEAT,"* it is essential to have the following correspondence:

Metaphorical Correspondence: (HeatisAnger)
Heat proportionally corresponds to emotional anger states.

However, unlike Example 2, it is not possible to deal with each metaphorical manifestation separately. Instead, one cocoon must be nested within the other. Given the rules shown, it is possible to infer the connotation that John had a tendency to angrily resent the verdict by a chain of inference partially shown in the left half of Figure 3.

Given this connotation, it can be argued that Example 3 indirectly implies that another tendency of John is not to angrily resent the verdict. This could be done by a scalar implicature (Hirschberg, 1985) outside the metaphor-pretense cocoon just as if the literal sentence *John had a tendency to angrily resent the verdict* had been uttered instead of the metaphorical one. The reasoning here would not be part of the metaphorical analysis of the sentence. However, another more metaphorical route is as follows: Some general pragmatic implicature is required specifying that when "one" person is mentioned in discourse, then it is reasonable to assume that there is at least one other person present, differing in a salient way from the one mentioned. For the purposes of this article we simplify by using the following defeasible rule:

SeveralPeople:
When one person in a group is explicitly mentioned then there is another person in the group, lacking the mental properties of the first person.

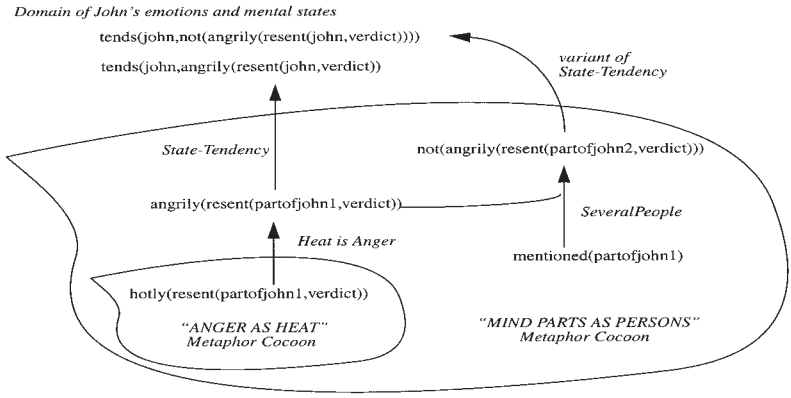

FIGURE 3 Dealing with the serial mixed metaphor *"One part of John hotly resented the verdict."*

Then ATT-Meta can infer within the outer cocoon that there is another subperson of John that does not angrily resent the verdict. From this, a negative variant of State-Tendency can create the inference that John also has tendencies not to angrily resent the verdict. This is sketched on the right-hand side of Figure 3.

The amount of reasoning within cocoons in the examples in this article is quite small and comparable to the extent of within-vehicle reasoning performed by, say, the MIDAS system (Martin, 1990). However, that was partly because of choice of example and partly because of deliberate simplification of the examples for purpose of study. In general, our approach countenances much more elaborate reasoning within cocoons, and this is reflected in the complexity of reasoning involved in other examples we have run ATT-Meta on (see, e.g., Barnden & Lee, 1999).

TOWARD PSYCHOLOGICAL PREDICTIONS

Generating psychological predictions from our model, even in nonmixed cases of metaphor, has not been a major focus of our work to date. However, because in the future some version of ATT-Meta could be put forward as a psychological model, we have in fact been concerned that ATT-Meta should, if possible, be broadly consistent with experimental results in psychology. In particular, we are sensitive to the debate in the psychological literature on metaphor about the relative amounts of time taken by people to understand metaphorical and literal utterances. Brisard, Frisson, and Sandra (2001/this issue) provide one example of work within this debate.

Our approach requires some sort of vehicle-domain meaning (roughly, literal meaning) to be constructed as a base for inference inside the pretense cocoon. It is

therefore, in some sense, a "literal first" approach, and some experimental findings have been put forward as conflicting with such approaches (see the literature review in Brisard et al., 2001/this issue). However, note two very important things. First, our approach is not of the type that says that metaphorical connotations are only sought once a literal meaning is discarded. As far as our approach is concerned, a metaphorical utterance could well have one or more literal meanings that make sense in context. Second, measures of whole-utterance reading times or understanding times may well be including a stretch of time for the understander to fully link the utterance to the surrounding discourse and to elaborate the basic compositional meaning of the utterance to get a full communicated proposition. By contrast, the sort of basic literal meaning that goes into the pretense cocoon in our approach does not need to be fully linked or elaborated in the way it would if it were taken to be the point of the utterance. (It may only need certain basic types of linkage, involving, for example, the choice of alternative literal senses for words based on sense choices already made in surrounding discourse.) Rather, it is precisely the within-cocoon reasoning steps and conversion rule applications that serve to link the literal meaning to the meaning of surrounding discourse. Thus, the within-cocoon reasoning and conversion rule applications *replace* discourse linking work that would normally have been done for the literal meaning had it been taken as the point of the utterance. After all, the vehicle-domain inferences that are needed within the cocoon, or inferences like them, could well be needed had the utterance been taken literally in a different context. It may be, therefore, that if our model has an overhead it will lie mainly in the conversion rule applications. However, given that these applications in general only account for a minority of the overall inference steps, the overhead may be minor.

In sum, it is simplistic to assume that a model that espouses the generation of some sort of literal meaning as a basis for the generation of metaphorical meaning, as our model does, has any implication that the processing time needed for the utterance is equal to the time needed to process the utterance *had it been* taken literally plus extra time for doing such operations as metaphorical mapping (e.g., conversion rule application in our terms). It is just such simplistic assumptions that seem to lurk in much writing surrounding psychological experiments on metaphor. Basically, the problem is a simplistic view of what literal meaning is and how it is produced, forgetting in particular the needed discourse linking effort. Honeck, Welge, and Temple (1998) also noted that such effort is usually not discussed.

Having said all this, it is clear that the more it can be shown that metaphorical understanding does often take longer than literal understanding, the less pressure there is on us to put forward these arguments. Here we are encouraged by the results of Brisard et al. (2001/this issue) and others, who do find a slowdown during metaphorical processing. Also, we should note that the psychological results on processing time tend to focus on *"A is B"* style metaphorical utterances. Not only do such utterances form a minority of metaphorical utterances in mundane dis-

course (as opposed to discussions of metaphor), but also there is reason to think that different forms of metaphorical utterance lead to different timing comparison results (see, e.g., Onishi & Murphy, 1993).

Because the processing in our model is goal directed (see Barnden, 1998; Barnden & Lee, 1999), where the goals are supposed to arise from the processing of context, it is to be expected that the more relevant context there is, the more quickly and easily a given metaphorical utterance embedded within it will be processed. This is broadly consistent with experimental results.

It is not clear that our model currently leads to useful predictions about the particular case of mixed metaphors, other than that the more metaphors that are involved in an utterance, the greater will be the metaphorical slowdown, such as it is. However, any model that did not definitely claim equal time for metaphorical and literal processing would probably predict this.

Finally, we take issue with authors who take equal-time results for figurative and literal understanding to imply that no special type of processing is going on the figurative case. The most that can be inferred from an equal-time result is that if there is special processing, it takes no longer than nonspecial processing does. In any case, what counts as a special type of processing? It is a highly relevant, purpose-sensitive question. For example, is the reasoning inside an ATT-Meta cocoon special? Are conversion rule applications special? The answer to both of these questions is positive, in the sense that a cocoon is involved, and negative to both, in that each individual step is just an inference step of exactly the same computational sort as would be used for literal understanding.

FURTHER DISCUSSION

It is clear that parallel mixes present fewer difficulties to any preexisting theory of metaphor than serial mixes. This is due to the frequent lack of interaction between the two metaphors involved. However, this is not to say that the metaphors in parallel mixing operate in total isolation. Certain parallel mixes are more common than others. In particular, metaphors that refer to abstract entities as physical objects are often mixed with spatial metaphors. For example:

4. *"John pushed the ideas to the back of his mind."*

Example 4 uses two familiar conceptual metaphors: *"IDEAS AS PHYSICAL OBJECTS"* and *"MIND AS ENCLOSED SPACE."* However, it is not clear whether such examples are instances of live mixing. There are two reasons for doubt. First, such examples can often be termed *dead mixes*, mixes that have been so conventionalized that there is no need for any extra reasoning to combine the two familiar metaphors. This, however, is not to suggest that the individual metaphors are dead, only that

the mix is so familiar that any metaphorical reasoning is performed in just one metaphor-pretense cocoon that represents the conventionalized mix of the metaphors.

Second, it is not clear whether the level of representation of conceptual metaphors is universal. It is conceivable that two different native speakers may represent the same metaphor with different levels of granularity, and, in some cases, a manifestation might be mixed to one speaker and straight to another. Therefore, to avoid such issues, we have adopted a position of methodological solipsism (Fodor, 1980) with respect to the particular set of metaphorical concepts assumed and focused on the actual processes and types of data structures involved in reasoning. Grady (1997) argued that certain metaphors heretofore considered as unitary (including "*THEORIES AS BUILDINGS*," in fact) should be regarded as mixes of finer grain metaphors. We agree with this, but there is still the question of which metaphors should be so viewed and how live or dead any mixing is in a given case.

In our brief references to parallel mixing in earlier work (e.g., Barnden, 1997a), we suggested that standard mixes can be handled by having a single metaphorical-pretense cocoon, instead of the two assumed in this article. That is, we have previously taken the one-cocoon approach as the default. In this approach, information in the two vehicle domains can interact. This could be seen as a form of blending (Turner & Fauconnier, 1995) with the pretense cocoon acting as the blend space. Sometimes such interaction is benign and easy to perform, and sometimes it is fraught with conflict (as in Examples 1 and 2). It is a matter of further research to reconcile the one-cocoon and two-cocoon approaches. One factor that is involved may be the extent of the cognitive resources available: Because the one-cocoon approach may have to deal with conflicts between the two metaphors, it should perhaps only be attempted (in cases of live mixing) when there are cognitive resources to spare.

In serial mixes, the metaphors strongly interact. If the analysis provided earlier is correct, and serial metaphors work by the chaining of one vehicle domain to the other vehicle domain to the tenor, then conversion rules provide an explicit constraint on what metaphors can be sensibly mixed because a sensible mapping is required from the former vehicle to the latter.

In this view, conversion rules act as filters between domains, first, to constrain the types of serial mixed metaphor possible, and second, to constrain the types of information transferred as only metaphorical manifestations that make sense in the other metaphor-pretense cocoon can be mapped.

In our previous work, it has been assumed that generality in conversion rules and mapping is a good thing. However, given this filtering role, specificity is an advantage because it provides strong constraints on mixing. Clearly, within-vehicle reasoning is important here. If more specific conversion rules are favored, then more of the reasoning workload must be performed prior to mapping to the tenor domain.

In this article, we have dealt exclusively with the mixing of metaphors. However, the mixing of different tropes is also common. For example, Warren (1992) classified five possible combinations of metonymy and metaphor: metaphor within meto-

nyms, metonyms within metaphors, metaphors within metonyms, metaphor in metaphor, and metonyms in metonyms. The only other relevant computational work we are aware of is Fass's (1997) meta* system, which is capable of understanding metonymy and metaphor mixes. (Fass, 1997, also addressed serially mixed metaphor.) In collecting data we have ourselves noted a particular type of metonymy that often occurs in combination with metaphors of mind. An example is *"China was at the surface of John's mind."* Because it is presumably some idea of China, not China itself, that is in the physical space suggested by the "surface of" wording, we have here a combination of a *"THING FOR IDEA OF IT"* metonymy with a *"MIND AS PHYSICAL SPACE"* metaphor. Another interesting example of potential metonymy and metaphor mixing is *"Sally tore Mike's talk to shreds,"* which could variously be interpreted as involving just a metonymy going from the talk to the paper on which the talk was written, so that the tearing is literal, or as involving no metonymy but instead a metaphor of a talk event as a piece of physical fabric, or finally as involving a metonymic link from the talk to the ideas in the talk combined with a metaphor of a body of ideas as a piece of physical fabric.

Also, D. Fass (personal communication, October, 1999) suggested an alternative analysis of Example 3, *"One part of John hotly resented the verdict,"* involving metonymy.[1] Under this analysis, there are two separate "part of" operations in the example: (a) a part of John that is John's emotional and mental states and (b) a part of John's emotional and mental states. Because the explicit mention of "part of" seems to refer to the second sense, the first must be expressed metonymically. Although metonymy is not currently implemented within ATT-Meta, there is no reason why it could not. One future research goal is an analysis of the interaction of metaphor with other tropes such as metonymy. Which particular interpretation is preferred depends on both the context of the metaphor and the particular conceptual knowledge of the hearer.

CONCLUSIONS

In this article, we described some initial work on mixed metaphors. We argued that both parallel and serial mixes can be processed using basic AI reasoning techniques that have already been applied to cases of unmixed metaphor. The serial case requires something extra: the nesting of metaphorical-pretense contexts (cocoons) within each other. However, as the cocoons are similar to the simulative reasoning cocoons also used in ATT-Meta to reason nonmetaphorically about agents' beliefs, and those cocoons also need to be mutually nested, the mutual nesting of metaphorical cocoons is not a conceptually major addition.

---

[1]Despite his analysis being technically possible and consistent with our approach, we find our analysis more plausible owing to its greater simplicity.

We have also suggested that within-vehicle reasoning plays an important role in unmixed metaphor, and this role extends naturally to mixed cases. Indeed, because the point of within-vehicle reasoning is to connect the vehicle-domain content of a metaphorical utterance to vehicle-domain concepts that the known mappings can directly handle (to avoid as far as possible the expensive process of discovering new mappings), within-vehicle reasoning plays a particularly important role in mixing because of the higher number of domains being juggled. Questions for further research include that of criteria for choosing to pursue a serial interpretation versus a parallel one, and that of criteria for deciding during processing in the parallel case whether to use one metaphorical-pretense cocoon or two.

ACKNOWLEDGMENTS

This research is being supported in part by Grant GR/M64208 from the Engineering and Physical Sciences Research Council, England, and has previously been supported by Grant IRI–9101354 from the National Science Foundation.

REFERENCES

Barnden, J. A. (1997a). An AI system for metaphorical reasoning about mental states in discourse. In J.-P. Koenig (Ed.), *Conceptual structure, discourse and language* (Vol. 2, pp. 167–188). Stanford, CA: Center for the Study of Language and Information.

Barnden, J. A. (1997b). Consciousness and common-sense metaphors of the mind. In S. O'Nuallain, P. McKevitt, & E. Mac Aogain (Eds.), *Two sciences of the mind: Readings in cognitive science and consciousness* (pp. 311–340). Amsterdam: Benjamins.

Barnden, J. A. (1998). Combining uncertain belief reasoning and uncertain metaphor-based reasoning. In M. A. Gernsbacher & S. J. Derry (Eds.), *Proceedings of the 20th Annual Meeting of the Cognitive Science Society* (pp. 114–119). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Barnden, J. A., Helmrich, S., Iverson, E., & Stein, G. C. (1996). Artificial intelligence and metaphors of mind: Within-vehicle reasoning and its benefits. *Metaphor and Symbolic Activity, 11,* 101–123.

Barnden, J. A., & Lee, M. G. (1999). An implemented context system that combines belief reasoning, metaphor-based reasoning and uncertainty handling. In P. Bouquet, P. Brezillon, & L. Serafini (Eds.), *Modelling and using context* (Lecture Notes in Artificial Intelligence, No. 1688, pp. 28–41). Berlin, Germany: Springer-Verlag.

Brisard, F., Frisson, S., & Sandra, D. (2001/this issue). Processing unfamiliar metaphors in a self-paced reading task. *Metaphor and Symbol, 16,* 87–108.

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

Fodor, J. A. (1980). Methodological solipsism considered as a research strategy in cognitive psychology. *Behavioural and Brain Sciences, 3,* 63–109.

Fowler, H. W. (1908). *The king's English* (2nd ed.). Oxford, England: Clarendon.

Grady, J. E. (1997). Theories are buildings revisited. *Cognitive Linguistics, 8,* 267–290.

Hirschberg, J. (1985). *A theory of scalar implicature.* Unpublished doctoral dissertation, University of Pennsylvania, Philadelphia.

Honeck, R. P., Welge, J., & Temple, J. G. (1998). The symmetry control in tests of the standard pragmatic model: The case of proverb comprehension. *Metaphor and Symbol, 13,* 257–273.

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

Lakoff, G. (1996). Sorry, I'm not myself today: The metaphor system for conceptualizing the self. In G. Fauconnier & E. Sweetser (Eds.), *Spaces, worlds, and grammar* (pp. 91–123). Chicago: University of Chicago Press.

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

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

Onishi, K. H., & Murphy, G. L. (1993). Metaphoric reference: When metaphors are not understood as easily as literal expressions. *Memory and Cognition, 21,* 763–772.

Turner, M., & Fauconnier, G. (1995). Conceptual integration and formal expression. *Metaphor and Symbolic Activity, 10,* 183–204.

Warren, B. (1992). *Sense developments: A contrastive study of the development of slang senses and novel standard senses in English.* Stockholm: Almqvist & Wiksell.