

## METAPHOR IN THE MIND<sup>1</sup>

### §1: GENERAL ISSUES IN THE STUDY OF METAPHOR

Philosophers have often adopted a dismissive attitude toward metaphor. Hobbes (1651, ch. 8) advocated excluding metaphors from rational discourse because they “openly profess deceit,” while Locke (1690, Bk. 3, ch. 10) claimed that figurative uses of language serve only “to insinuate wrong ideas, move the passions, and thereby mislead the judgment; and so indeed are perfect cheats.” Later, logical positivists like Ayer and Carnap assumed that because metaphors like

(1) “How sweet the moonlight sleeps upon this bank!”<sup>2</sup>

involve category mistakes, they can have no real meaning or verification conditions. Thus, they also mentioned metaphor only to place it beyond the pale of rational discourse.

Starting in the 1960’s and 70’s, philosophers and linguists began to take more positive interest in metaphor. In particular, they proposed that the literal absurdity of metaphors like (1) induces a “metaphorical twist” (Beardsley 1962) which endows the sentence with a new semantic meaning. Some, like Beardsley, suggested that the “twist” converts the metaphorical expression’s normal connotation into its denotation, while others, like Weinreich (1966) and Levin (1977) argued that it alters the selection restrictions and semantic markers within the expression’s lexical entry (that is, the specification of what sorts of objects the expression can apply to and what sorts of other expressions it can combine with). These theories took metaphors seriously; but they suffered from two basic flaws. First, not all metaphors are literally false or absurd (e.g. Cohen 1975): consider

(2) No man is an island.

(3) Anchorage is a cold city.

Such cases thus need some alternative mechanism by which their meaning can be “twisted.” More fundamentally, these theories all assumed that the “twist” could only operate on features that were already somehow part of words’ literal meanings. But nearly any word can be used metaphorically, and any

---

<sup>1</sup> Elisabeth Camp; Harvard Society of Fellows; [ecamp@fas.harvard.edu](mailto:ecamp@fas.harvard.edu).

<sup>2</sup> *The Merchant of Venice*, V.i.54.

given word can produce a wide range of metaphorical meanings in different contexts. Including all the materials needed to produce all these meanings within the lexicon effectively exploded the semantic theory altogether.

In retreat from this semantic deluge, theorists turned either to *pragmatic* theories (Grice 1975, Searle 1979, Martinich 1984), on which the speaker says (or “makes as if to say”) a sentence with its normal, literal meaning and thereby communicates a distinct propositional content; or else to so-called *non-cognitivist* theories (Davidson 1978), which reject the notion of metaphorical meaning altogether and focus on the non-propositional shifts of “perspective” that metaphors can induce.

Recently, several philosophers have argued that these views fail to do justice to the role that metaphorical utterances actually play in conversation. We regularly use metaphors to make assertions and other speech acts with more or less determinate contents, but the non-cognitivist is committed to denying this. Further, within a given conversational context, a metaphorical interpretation may become the default way of using certain words. If pragmatist is committed to treating metaphor as a form of particularized conversational implicature, as Grice suggests, then this fact is difficult to explain. Josef Stern (2000) accommodates these facts by applying semantic techniques for capturing the context-sensitivity of expressions like “I” and “that” to metaphor. By contrast, theorists in the tradition of Relevance theory (e.g. Sperber and Wilson (1995, 1986), Bezuidenhout (2001)) and other “contextualists” like Recanati (2004, 2001) advocate assimilating metaphor to the broad category of “pragmatic intrusion” into “what is said.” (Pragmatic intrusion involves both “enrichment,” as when my utterance of

(4) I’ve had breakfast.

communicates that I’ve eaten breakfast *today*, and “loose talk,” as when I use “silent” to describe a house that is quiet except for the humming refrigerator and dripping faucet.) I am skeptical whether either a formalist or contextualist analysis of metaphor is really warranted (Camp 2005, 2006). But these treatments have had the quite salutary effect of focusing attention within philosophy on the specific ways in which metaphor actually works.

In particular, contextualists have emphasized that metaphor is not just an occasional rhetorical flourish, but a pervasive and often comparatively unreflective aspect of ordinary language use. This fact prompts several questions: What is it about our minds that makes metaphor so natural? Is it something distinctive to metaphor, or a general feature of cognition? Indeed, could metaphor be a fundamental feature of thought itself? Because the most sustained and innovative contemporary work on metaphor has come from empirical and theoretical research in psychology and cognitive science, and because this work is relatively unfamiliar to philosophers, I'll focus on it here, noting connections to discussions within philosophy where appropriate. In the remainder of this section, I'll take up the question of whether metaphor is processed differently from literal speech. In §2, I discuss the four main models of metaphor comprehension.

Within psycholinguistics, the main topic of investigation has been whether metaphorical comprehension is “direct” or “indirect.” On the “direct” model, the metaphorical meaning is accessed immediately, while the “indirect” model postulates that the hearer seeks a metaphorical interpretation only after the search for a plausible literal meaning fails. Investigators have tended to assume that answering the question of directness should help adjudicate between a pragmatic Gricean analysis and a contextualist (or semantic) one. It should be noted that both Grice and Searle intended their theories as rational reconstructions, aimed at explaining how utterances could in principle enable successful communication. They were not concerned with how utterances are actually processed, let alone with the conscious experience of linguistic interpretation. Nonetheless, their views do naturally suggest a certain processing model, and the question of comprehension is an independently interesting one.

A range of studies support the direct model, and until relatively recently it has been the dominant view. Several studies by Gibbs (e.g. 1990, 1994) found no difference between the time it took hearers to comprehend literal and metaphorical speech, and specifically to comprehend literal and metaphorical anaphoric descriptions, as in

- (5) There was one boxer that Stu hated. This guy always lost....{*The creampuff/ The loser*} didn't even show up (Gibbs 1990, 59).

In another important study, Glucksberg, Gildea, and Bookin (1982) found that it took subjects longer to decide whether a sentence was literally true when the sentence also had a plausible metaphorical meaning, even though the subjects had been instructed to consider only the literal truth-value. This suggests that metaphorical interpretation arises automatically, and not merely after a failure of literal interpretation.

More recently, however, these findings have been significantly qualified. Various studies (e.g. Blasko and Connine 1993, Gentner and Wolff 1997, Bowdle and Gentner 2005, Giora 1997, 2002) have found that unfamiliar and novel metaphors take significantly longer to process than either literal sentences or familiar metaphors. Bowdle and Gentner (2005) also found that novel similes are processed significantly faster than novel metaphors, suggesting that it's not merely the unfamiliar juxtaposition of terms, but the literal sentence meaning itself, that increases processing time. In general, while the surrounding context of utterance can influence comprehension time (Ortony et al 1978), the most influential factors appear to be the broader familiarity or prototypicality of that way of using those words (Giora 1997, 2002, 2003). There also appears to be an independent effect for a metaphor's *aptness*:<sup>3</sup> among unfamiliar metaphors, highly apt meanings are interpreted rapidly, although not as quickly as literal meanings, while merely moderately apt metaphors take significantly longer (Blasko and Connine 1993, Gentner and Wolff 1997).

In addition to studies of processing time, psycholinguists have also investigated the interpretation of metaphor and of irony in people with cognitive and brain disorders. The overall picture which emerges is that the comprehension of both metaphor and irony involves the "theory of mind" — the ability to attribute thoughts to others — in a way that literal meaning does not; in contrast to irony, however, metaphor requires only a first-order theory of mind, and not an ability to attribute thoughts about thoughts to others. Happé (1993, 1995) found this pattern clearly demonstrated with autistic subjects. Similarly, several investigations (e.g. Brownell et al 1990, Brownell et al 1984, Winner and Gardner 1977) have found that subjects with right-hemisphere damage, but not those with left-hemisphere damage, are biased

---

<sup>3</sup> Aptness was determined by asking subjects to rate "how well you think the metaphor expresses it's [sic] specific non-literal meaning" (Blasko and Connine 1993, 297).

toward literal interpretation. They can interpret highly conventional metaphors but have difficulty selecting which meaning, literal or metaphorical, is contextually appropriate, and have more difficulty with metaphorical polysemous alternatives than with nonmetaphorical ones. (Very roughly, distinctively linguistic processing tends to be concentrated in the left hemisphere, while the right hemisphere is concerned with broader cognitive tasks, including pragmatic interpretation and emotive and social response.)<sup>4</sup>

Overall, then, the empirical data clearly militate against drawing a sharp boundary between literal and metaphorical meaning, or insisting that metaphor is a “deviant” or unusual use of language. But at the same time, they do clearly support some distinction between literal and metaphorical meaning.

As we saw, advocates of the “direct” model take the pervasiveness and unreflectiveness of our metaphorical speech to show that metaphors are processed in the same way as literal utterances. By contrast, “cognitive linguists” in the tradition of George Lakoff take these same facts to show that even apparently literal utterances are processed metaphorically: that is, that “our ordinary conceptual system, in terms of which we both think and act, is fundamentally metaphorical in nature” (Lakoff and Johnson 1980, 4). To establish this claim, they point to the fact that we so often and so naturally say things like

(6) Your claims are *indefensible*.

(7) He *attacked* every weak point in my argument.

(8) I *demolished* his argument.

Statements like these seem to reveal an underlying, unreflective metaphorical conceptualization: in this case, of arguments in terms of war. As Lakoff says, “The words and fixed expressions of a language can

---

<sup>4</sup> The findings with schizophrenics are less conclusive. Billow et al (1997) analyzed psychiatric interviews with schizophrenics and with cognitively normal medical patients, and found no difference in the overall incidence of metaphor, although the schizophrenics’ utterances and interpretations were more often “autistic” and “tangential” than for controls’. By contrast, Langdon et al (2002) did find significant, and independent, impairments for interpreting irony and metaphor in schizophrenics, even after controlling for the inability to inhibit prepotent information; thus, the obstacle to comprehension seems not to be just an inability to ignore literal meaning. However, they also found that theory of mind made no independent contribution to the ability to comprehend metaphor (as opposed to irony); the authors suggest that the obstacle to comprehending metaphor may be schizophrenics’ disorganized “semantic networks.”

*code*, that is, be used to express aspects of, a given conceptual metaphor” (1987, 384, emphasis in original). Similarly, instances of polysemy — for instance, the many meanings of “over,” as in

(9) The plane is flying *over* the hill.

(10) Sam turned the page *over*.

(11) The play is *over*.

— often seem to result from metaphorical extrapolation from a basic, concrete sense to more abstract applications (1987, 419).

Once again, these facts about our normal patterns of speech are notable. However, the studies cited above show that one must be careful about relying on highly conventionalized metaphors to establish general conclusions about metaphor, language, or thought. As metaphors become lexicalized, they are no longer processed *as* metaphors, even if we believe we can discern their metaphorical roots. Metaphor is an important mechanism for language change. But the fact that many ways of using words seem to have originally been metaphorical doesn’t itself show that the thoughts we now use them to express involve metaphorical cross-domain mappings.<sup>5</sup>

Still, even putting such highly conventional examples aside, the fact remains that we find it very natural to employ certain metaphorical tropes for talking about certain topics, such as arguments or love. This fact does call for explanation, and it’s natural to seek an explanation for it within cognition itself. Lakoff’s explanation is that our thoughts about these topics are necessarily metaphorical, because the topics cannot “be fully comprehended on their own terms. Instead, we must understand them in terms of other entities and experiences” (Lakoff and Johnson 1980, 177). The claim that metaphor is unavoidable may well be true for thought about unfamiliar or mysterious topics, such as computers or particle physics. But even here, this doesn’t imply that we can’t think about — that is, refer in thought to — those topics at

---

<sup>5</sup> It’s not even obvious that those uses ever were genuinely metaphorical: our intuitions about words’ ancestries are not generally very reliable. Keysar and Bly (1995) found that subjects who were presented with an unfamiliar idiom in one of two contexts — for instance, “The goose hangs high” in the context of a sad or of a happy story — not only attributed meanings to the idiom that comported with the story, but were sure that they would have attributed the same meaning if they had been presented with the idiom even without the context; they also believed that it was highly unlikely that the idiom could have the alternate meaning.

all except in metaphorical terms. Metaphors are powerful tools for structuring thought, and they do make it easy to focus on certain facts and ignore others. But with the possible exception of extremely abstract subjects, like God, we can think about the relevant topics directly enough to discover facts that the metaphorical mapping can't accommodate. Idealist arguments aside, no one should grant that metaphors impose an absolute filter on which information we are capable of cognizing.

Lakoff and his colleagues aren't primarily concerned with thought about unfamiliar topics, though; they're most interested in establishing the metaphorical nature of ordinary thought about familiar matters like arguments or anger. They maintain that we cannot think about these topics "in their own terms" either, because they are also too abstract; this leaves us no choice but to filter them through our embodied experience with more concrete domains. The class of metaphors for which this hypothesis is most compelling is the spatial representation of more abstract domains. (In particular, subjects' estimates about the passage of time do seem both to depend upon the particular spatial metaphors for time prevalent in their native language, and to be affected by non-verbal spatial information; but there is no inverse dependence of spatial representations upon temporal metaphors (Casasanto 2005).) However, for many of Lakoff's paradigm cases, it's highly unlikely that the relevant asymmetry in direct cognizability obtains: our experiences of these topics are at least as embodied and concrete, and are accessible at least as early in life, as our experiences of the domains in whose terms we characterize them metaphorically. For instance, Lakoff often cites the "conceptual metaphor" ANGER IS HEAT OF FLUID IN A CONTAINER (e.g. 1987, 384-388). But children experience anger well before they understand the effects of heat on fluid pressure in closed containers (Ortony 1988). Similarly, I've been in many arguments, but I have only very little, very indirect experience with war; and that experience is itself quite unconnected to the highly strategic aspects of war that underwrite metaphorical descriptions of arguments.

*Prima facie*, a much weaker hypothesis suffices to explain much of the data to which Lakoff and his colleagues rightly bring our attention. We do find metaphorical talk natural and useful because of the role that metaphors play within cognition; but this need not imply that our representation of the one domain is entirely parasitic on our representation of the other. It might just be that metaphors reflect and

reinforce similarities between our independent mental representations of the two domains: for instance, our distinct representations of arguments and of war (cf. Murphy 1996, 179; Fogelin 1988, 86). A similarity-based view is compatible with a range of models of just how the source domain is mapped on to the topic domain, as we'll see in the next section, while the strong dependence thesis can really only be accommodated on one of the models. The similarity-based view is also compatible with the claim that how we think about a given topic is affected by the metaphors we regularly hear and employ, so that our mental representations are not wholly antecedent to and independent of metaphorical talk. While this claim is plausible, it is not restricted to metaphor: proverbs, fictions, and religious and scientific doctrines can have the same effect.

The claim that thought itself is metaphorical can seem unavoidable if the alternative is (a caricature of) an old-fashioned theory of mind, which construes cognition as operating exclusively on strings in a propositional logic, and as employing exclusively operations like deduction and inductive probabilification. Philosophers and cognitive scientists have often employed such a model as a convenient idealization, and sometimes as more than that. But in the last thirty years, psychologists, cognitive scientists, and even philosophers have also increasingly paid attention phenomena that don't fit easily within this model: for example, our embodied knowledge of how to drive to the mall or to order food in a restaurant, or our tendency to treat certain members of a class (say, sparrows) as better exemplars of that class (birds) than others (say, penguins). In order to explain these phenomena, psychologists and cognitive scientists have postulated prototypes (e.g. Rosch 1978), scripts (e.g. Shank and Abelson 1977), schemas (e.g. Holyoak and Thagard 1995), and other mental representations (Barsalou 1983, 1993) that are more complex, more contextually malleable, and more intimately connected to action and experience than the traditional picture seems to allow.<sup>6</sup> They typically call these mental representations "concepts"; like many philosophers, I find this terminology misleading, because these mental representations don't do much of the work that concepts have traditionally been assumed to

---

<sup>6</sup> Whether the traditional picture itself can accommodate all of these phenomena, or whether the difference between the two models is merely a matter of emphasis, is a further, open question.

do. (Most notably, they don't compose into whole thoughts such that the truth-conditions of the whole thought is a systematic function of the meanings of its parts.) For clarity, I'll call all these 'alternative' mental representations *characterizations*. Whatever we call them, though, these representations are fairly clearly not themselves metaphorical.

## §2: MODELS OF METAPHORICAL COMPREHENSION

How, then, do we understand metaphorical utterances? What effects do metaphors have in cognition? In many ways, this represents the core of the question of metaphor. Besides being an interesting topic in its own right, the answer may have important ramifications for psychological and philosophical theories of the mind more generally. Even if metaphorical utterances are processed differently than literal ones, the pervasiveness and naturalness of metaphorical speech does suggest that the sorts of mental representations and operations we employ in comprehending it are quite common and relatively efficient.

The answer to the question of how metaphors are comprehended is, however, largely independent of the taxonomic issue of metaphorical meaning discussed at the beginning of §1. Someone like Davidson rejects the notion of metaphorical meaning altogether: he points to "seeing as" as a non-propositional cognitive effect associated with metaphor, and stops there. But pragmatists like Grice and Searle, or contextualists like Recanati and Bezuidenhout, can appeal to those same effects as the mechanism by which hearers determine the speaker's meaning or the content of 'what is said.' Even a semantic theorist like Stern could invoke them as the means by which the contextually-sensitive semantic values of metaphorical expressions are fixed. Depending on one's view about the status of metaphorical meaning and which particular model of metaphorical effects one adopts, one will have more or less work to do after developing an account of metaphor's effects in order to derive appropriate metaphorical meanings. But barring substantial additional assumptions, the theoretical relationship between actual cognitive processes and metaphor's linguistic status remains open.

Before turning to the four main types of models that have been developed, we should flag a general complication that is not much noticed among cognitive scientists, nor always by philosophers. Metaphors come in all syntactic shapes and sizes. They aren't restricted to 'noun-noun' constructions in which the subject is literal and the predicate is metaphorical, as in

(12) Juliet is the sun.

Rather, the noun phrase itself can be metaphorical, as in

(13) The fox is fomenting discord among us once more.

The predicate can be an active verb, as in

(14) The earth pirouettes around the sun.

Or the entire sentence can metaphorically describe an unmentioned situation, as in

(15) The sun blazes bright today; the clouds flee from his mighty beams,

describing Achilles as he rages upon the battlefield. All of these examples introduce interpretive complexity. In each case, however, there is still some topic, even if it is only implicitly identified; and this topic is supposed to be thought of in terms of something else, which I'll call the 'source'. All four models aim to explain how these two domains — source and topic — interact.

On the simplest model of metaphorical comprehension (cf. Rorty 1987), a metaphor merely juxtaposes the topic under discussion (e.g. Juliet) with another object, event or situation (e.g. the sun), and thereby causes us to notice surprising features of the topic. While this view is appealingly minimal, it suffers from the fundamental flaw of being too accommodating. Hearing an utterance of (12) might remind me of a particularly pleasant sunny day at the beach, which might in turn lead me to notice that Juliet's eyes are the color of the sea on that day. But even apart from the question of metaphorical *meaning*, this effect is clearly irrelevant: it is a mere idiosyncratic association, not appropriately related to the project of thinking of Juliet *as the sun*. A second problem is that the juxtaposition view cannot explain how metaphors manage to do more than merely "nudge us into noting" already known but neglected features (Davidson 1985, 253). For instance, hearing someone assert (15) doesn't just make me reconfigure my thoughts about Achilles in a certain way, as Davidson claims. It also informs me of

something new: that Achilles is fighting with great energy and force. Because a juxtaposition view is limited to effects generated by juxtaposing my existing characterizations of the topic and source, it cannot account for this informativeness.

“Category-transfer” models solve the first problem by limiting the range of relevant features to those that are grounded in our characterization of the source. For instance, Nelson Goodman (1968, 72) claims that in metaphor, “a label along with others constituting a schema is in effect detached from the home realm of that schema and applied for the sorting and organizing of an alien realm.” Goodman’s brief sketch leaves mysterious just how a schema can organize an “alien” realm to which it cannot literally apply: he says that “a metaphor is an affair between a predicate with a past and an object that yields while protesting” (1968, 69), but he offers no specification of what “yielding” and “protesting” amount to. More recently, Glucksberg and Keysar (1990, 1993) have offered a cognitive model that dispenses with Goodman’s resolute nominalism; on their view, metaphorical comprehension and thought involves forming an *ad hoc* category from the metaphorical source by abstraction from a prototypical instance of the literally denoted category. (This is the model of metaphorical comprehension most amenable to the Lakoffian claim that our representation of the target domain is fundamentally dependent on our representation of the source domain.) So, for instance, when we think of a job as a jail, we abstract away from specific, concrete features of jails to produce a general schema which includes being involuntary, unpleasant, confining, punishing, unrewarding, and so on (Glucksberg and Keysar 1990, 7). The category-transfer model also has a solution to the problem of informativeness: by classifying the topic within the generated category, the metaphor prompts us to add any missing features of that category to our characterization of the topic.

Because category-transfer models focus on the whole, complex schema associated with the source, they nicely explain the global organizational effect that is such a prominent part of metaphorical comprehension. As Goodman says, in metaphor not just an isolated term, but “a whole apparatus of organization, takes over new territory” (1968, 73); Glucksberg and Keysar (1993, 421) say that metaphors present “a patterned complex of properties in one chunk.” However, because these models focus

exclusively on the schema associated with the *source*, they have a hard time explaining the very different effects that can be produced by applying the same vehicle to different topics. When we think of Juliet as the sun, for instance, the result is very different than when we think of Achilles as the sun. A natural fix (cf. Glucksberg et al 1997) is to assign the topic a filtering or selectional role: only those aspects of the schema are transferred which capture dimensions along which one might relevantly and significantly classify that topic. This helps; but still, precisely because the topic can only filter *out* features that are already generated by abstraction from the source, the model faces the same problem that plagued lexical semantic theories sketched at the outset: the schema must be sufficiently general and encompassing to apply to any possible topic. For instance, the abstracted schema for *the sun* must include all the relevant features to apply, not merely to Juliet and Achilles, but also to Louis XIV, Richard III, God, an atomic bomb, the nucleus, the Platonic Forms, and so on. This difficulty is significantly mitigated by making the generated schema sensitive to the current (cognitive or communicative) context (cf. Glucksberg et al 1997, Stern 2000). But even this doesn't eliminate the problem altogether, because the same metaphorical source can be applied to very different topics within a single context, as in

(16) Juliet, Achilles, and Louis XIV are all suns in their own ways.

Finally, because the schema is only generated by abstraction, the model has a difficult time explaining the very concrete, experiential properties that metaphors can invoke (cf. Camp forthcoming). For instance, my job might be like a jail in the specific respects of requiring me to share a small cubicle with someone I don't like and to eat tasteless food dished out by surly staff. These features of the topic are also features of the source, and they well might be part of what I mean when I say

(17) My job is a jail.

But if the model includes these features in the schema, as it must in order to account for my use of (17), then it should be committed to predicting that they will also apply when someone says

(18) My marriage as a jail.

These are also possible features for that topic, and so they shouldn't be filtered out in principle. But in a relevantly similar pair of contexts in which hearers are equally ignorant about the speaker's job and

marriage, the features of sharing a small cubicle and eating tasteless food might well be included in the metaphorical effects of (17) but not of (18).

Rather than locating all the action exclusively in the source and assigning the topic at most a filtering role, “feature-matching” models (e.g. Ortony 1979, Fogelin 1988) operate through direct comparison with the topic. As a result, they nicely explain why the same source has such different effects when applied to distinct topics, and how metaphors can lead us to us notice very specific features of the topic. They also account for the traditional intuition that metaphors crucially involve comparison — that, as Quintilian said, “a metaphor is a brief similitude contracted into a single word.” Modern feature-matching models avoid the criticism that similarity is vacuous (cf. Goodman 1972) by exploiting Tversky’s (1977) salience-based theory of similarity. The extent to which two things count as similar in a given context, on this view, is determined by a weighted function of their shared salient features minus a weighted function of their distinctive salient features; a given feature’s *salience* is in turn understood as a function of its *intensity* and its *diagnosticity* in that context.<sup>7</sup> What distinguishes figurative comparisons (both metaphors and similes) from literal statements of similarity is that in the former but not the latter case, the source possesses highly salient features that cannot find matches in the topic, and the matches that can be established involve features that are highly salient in the source but not in the topic.<sup>8</sup> This dramatic *salience imbalance* is supposed to explain figurative comparisons’ irreversibility: why, for instance, it makes sense to think of a sermon as a sleeping pill, but not a sleeping pill as a sermon. Even where the ‘reversed’ metaphor is not anomalous, as with

(19a) That surgeon is a butcher.

(19b) That butcher is a surgeon.

---

<sup>7</sup> “Diagnosticity” refers to how useful the feature is for classifying objects in that context. The red stripes on a snake’s back might be quite *intense*, while the shape of its head might be more *diagnostic* of whether it is poisonous.

<sup>8</sup> Fogelin (1988, 89) suggests that the difference between figurative falsity and anomalousness is that in the former case, objects of the same general kind as the topic do possess features which match highly salient features in the source, while in the latter case even this is not possible.

differences in salience can still generate dramatically different bases of comparison in each case.<sup>9</sup>

The feature-matching model faces three main challenges. First, because it focuses on the search for matches between individual features, it has a hard time accounting for metaphors' holistic organizational effects. Second, because it searches for matches between already existing features in each characterization, it has a hard time explaining how metaphors manage to be informative. This difficulty is exacerbated by the fact that in figurative comparisons, many highly salient features in the source go unmatched: how do we distinguish those that we should ignore from those that we should introduce into the topic? Third, not all of the features that a metaphor intuitively makes us notice in the topic can be matched directly to features in the source. For instance, Romeo surely intends his metaphor as, in part, a paean to Juliet's beauty, but the sun is not itself beautiful.

The "structural alignment" view (Gentner 1983, Gentner and Wolff 1997, Gentner et al 2001, Bowdle and Gentner 2005) represents a kind of hybrid model of metaphor comprehension: it preserves the emphasis of category-transfer models on overall schemas and structural organization, but implements it within a comparativist framework. As a result, it can deploy the advantages of each view against the weaknesses of the other. The model postulates two stages of comprehension: alignment and projection. Within alignment, the processor begins by searching for matches in the topic to salient features in the source. Crucially, this includes indirect matches between identical higher-order relations that take a distinct argument in each characterization: for instance, *carry(blood)* would be matched to *carry(water)* in a comparison between blood vessels and aqueducts; this addresses the worry about the lack of direct matches. Local matches are then "coalesced" into "kernels" of internally consistent structures of related matches; the kernels are in turn merged into the largest overall coherent structure which preserves the greatest number of the largest kernels. The result is a complex, structured schema that is tailor-made for the particular topic. At the projection stage, features that are missing from the characterization of the topic but compatible with what is known about it, and whose presence would complete the structural

---

<sup>9</sup> But see Gentner and Clement (1998) and Tourangeo and Rips (1991) for evidence that salience imbalance is not predictive of figurativeness.

isomorphism of established matches, are added into the topic characterization; this solves the problem of informativeness.

Although the structural alignment view is more comprehensive than its rivals, it too is far from perfect. First, while metaphorical comprehension does crucially involve a global reorganization, this effect is often not nearly as systematic or consistent as the structural alignment model predicts; rather, it typically involves a more piecemeal patchwork of overlapping elements and kernels. Second, because the model places such heavy emphasis on higher-order structural matches, it has much the same difficulty as we saw with the category-transfer model in accounting for concrete features. Metaphors can attribute specific experiential properties to their topics which are neither direct projections from the source nor even projections from kernels in which all but one or two matches are established. For instance, if one describes a wine's taste as "velvet, with a brocade pattern,"<sup>10</sup> the primary communicated feature is a first-order, sensational property that the fabric itself does not possess. Such features may be produced by "blending" multiple source features (Fauconnier and Turner 1998), or they may "emerge" in a more mysterious fashion from the interaction between topic and source (Gineste et al 2000, Tourangeau and Rips 1991, Ricoeur 1977, Black 1962).

Ultimately, the three models are not necessarily as incompatible as their respective proponents often suggest. Ortony (1979) already includes both a recursive search for higher-order relational matches and the possibility of feature introduction within his feature-matching model, although Gentner's structural alignment model offers a more substantive explanation for just when these occur. The difference between category-transfer and structural-alignment models is more substantive, because they assign such different roles to the topic characterization. In fact, though, each model may be appropriate for a different sub-class of metaphors: structural alignment may capture the comprehension of novel metaphors, while category-transfer may be appropriate for more conventionalized ones (Gentner and Wolff 1997, Bowdle and Gentner 2005). This should not be surprising if we recall the evidence cited above for "direct" processing of established metaphors and "indirect" processing of novel ones. Highly

---

<sup>10</sup> I owe this example to Barry Smith.

familiar metaphors are on their way to becoming concepts in their own right; and as metaphors become conventionalized, they exhibit much less variation in their application to different topics.

Models of metaphorical comprehension like those I've discussed are theoretically and empirically useful because they offer relatively specific algorithms against which we can test actual thinkers' intuitions about particular cases. But precisely because they are algorithmic models, they are unlikely to replicate the full range of our intuitive responses to metaphors. Metaphor potentially involves the most creative aspects of human imagination and cognition. We should not expect it to be easy — to say the least — to reliably replicate leaps of imagination algorithmically. If anything, we should be surprised instead at how much interpersonal agreement and intuitive constraint there actually is in particular cases: at how little metaphorical interpretation is like mere idiosyncratic association.

It is fairly clear what puzzles and opportunities the comprehension of metaphor currently poses for cognitive science and psychology. At a minimum, the fact that metaphor depends so essentially on complex, nuanced mental representations, which in turn involve messy phenomena like similarity and salience, means that an adequate model of metaphor is a long way off. But it also means that developments in a fairly wide range of areas may end up being relevant.

The lessons for philosophers are less clear. As I mentioned above, facts about the process of metaphorical comprehension do not themselves decide the question about metaphor in which philosophers have been most interested: the taxonomic status of metaphorical meaning. Perhaps the most obvious lessons for philosophers are negative: do not assume that metaphor is a merely marginal aspect of thought or talk, or that there is no distinction between literal and metaphorical meaning, or that a substantive theory of similarity is impossible, without making more specific criticisms of the existing empirical research. Beyond this, philosophers who are not especially interested in metaphor *per se* may still find interesting applications for the empirical and theoretical work I've discussed here. Some examples that I find particularly intriguing include differences in how non-semantic speaker's meaning and conversational implicatures are processed; how to understand modes of mental representation that are not wholly propositional; the role of imagination in reasoning more generally; and the notion of 'aptness'

in aesthetic evaluation and in communication.

## References

- Barsalou, Lawrence (1993): "Flexibility, Structure, And Linguistic Vagary In Concepts: Manifestations Of A Compositional System Of Perceptual Symbols," in A.C. Collins, S.E. Gathercole, and M.A. Conway (eds.), *Theories of memory* (Hillsdale, NJ: Lawrence Erlbaum Associates).
- (1983): "Ad Hoc Categories," *Memory and Cognition* 11, 211-227.
- Beardsley, Monroe (1962): "The Metaphorical Twist," *Philosophy and Phenomenological Research* 22:3, 293-307.
- Bezuidenhout, Anne (2001): "Metaphor and What is Said: A Defense of a Direct Expression View of Metaphor," *Midwest Studies in Philosophy* 25, 156-186.
- Billow, Richard, Jeffrey Rossman, Nona Lewis, Deberah Goldman, and Charles Raps (1997): "Observing Expressive and Deviant Language in Schizophrenia," *Metaphor and Symbol* 12:3, 205-216.
- Black, Max (1962): "Metaphor," in *Models and Metaphors* (Ithaca, NY: Cornell University Press).
- Blasko, Dawn and Cynthia Connine (1993): "Effects Of Familiarity And Aptness On Metaphor Processing," *Journal of Experimental Psychology: Learning, Memory and Cognition* 19, 295-308.
- Bowdle, Brian and Dedre Gentner (2005): "The Career of Metaphor," *Psychological Review* 112:1, 193-216.
- Brownell, Hiram, Tracy Simpson, Amy Bihrl, Heather Potter, and Howard Gardner (1990): "Appreciation of metaphoric alternative word meanings by left and right brain-damaged patients," *Neuropsychologia* 28:4, 375-83.
- Brownell, Hiram, Heather Potter, Diane Michelow, and Howard Gardner (1984): "Sensitivity to Lexical Denotation and Connotation in Brain-Damaged Patients: A Double Dissociation?" *Brain and Language* 22, 253-265.
- Camp, Elisabeth (2006): "Metaphor, Contextualism, and What is Said," *Mind and Language* 21.
- (2005): "Critical Notice of Josef Stern's *Metaphor in Context*" *Noûs* 39:4, 716-732.
- (forthcoming): "Metaphor and That Certain 'Je Ne Sais Quoi'" *Philosophical Studies*.
- Casasanto, Daniel (2005): *Perceptual Foundations of Abstract Thought*. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Cohen, Ted (1975): "Figurative Speech and Figurative Acts," *Journal of Philosophy* 71, 669-684.
- Davidson, Donald (1978): "What Metaphors Mean," in S. Sacks (ed.), *On Metaphor* (Chicago: University of Chicago Press).
- Fauconnier, Gilles, and Mark Turner (1998): "Conceptual integration networks," *Cognitive Science* 22, 133-187.
- Fogelin, Robert (1988): *Figuratively Speaking* (New Haven: Yale University Press).
- Gentner, Dedre (1983): "Structure-mapping: A Theoretical Framework for Analogy," *Cognitive Science* 7, 155-170.
- and Brian F. Bowdle, Phillip Wolff, and Consuelo Boronat (2001): "Metaphor is Like Analogy," in D. Gentner, K. Holyoak, and B. Kokinov (eds.), *The Analogical Mind: Perspectives from Cognitive Science* (Cambridge, MA: MIT Press).
- and Catherine Clement (1988): "Evidence for relational selectivity in the interpretation of analogy and metaphor," in G. H. Bower (ed.), *The psychology of learning and motivation: advances in research and theory* 22 (New York: Academic Press).
- and Phillip Wolff (1997): "Alignment in the Processing of Metaphor," *Journal of Memory and Language* 37, 331-355.
- Gibbs, Ray (1994): *The Poetics of Mind: Figurative Thought, Language, and Understanding* (Cambridge: Cambridge University Press).
- (1990): "Comprehending Figurative Referential Descriptions," *Journal of Experimental Psychology: Learning, Memory and Cognition* 16:1, 56-66.
- GINESTE, Marie-Dominique, Bipin Indurkha, and Veronique Scart (2000): "Emergence of Features in Metaphor Comprehension," *Metaphor and Symbol* 15:3, 117-135.

- Giora, Rachel (2003): *On Our Mind: Salience, Context, and Figurative Language* (Oxford: Oxford University Press).
- (2002): “Literal vs. Figurative language: Different or Equal?” *Journal of Pragmatics* 34, 487–506.
- (1997): “Understanding Figurative and Literal Language: The Graded Salience Hypothesis,” *Cognitive Linguistics* 8, 183-206.
- Glucksberg, Sam and Boaz Keysar (1993): “How Metaphors Work” in A. Ortony (ed.), *Metaphor and Thought*, 2<sup>nd</sup> ed. (Cambridge: Cambridge University Press.)
- (1990): “Understanding Metaphorical Comparisons: Beyond Similarity,” *Psychological Review* 97, 3-18.
- Glucksberg, Sam, Matthew McGlone, and Deanna Manfredi (1997): “Property Attribution in Metaphor Comprehension,” *Journal of Memory and Language* 36, 50-67.
- Glucksberg, Sam, Patricia Gildea, and Howard Bookin (1982): “On Understanding Nonliteral Speech: Can People Ignore Metaphors?” *Journal of Verbal Learning and Verbal Behavior* 21, 85-98.
- Goodman, Nelson (1972): “Seven Strictures On Similarity,” in *Problems and Projects* (New York: Bobbs-Merrill).
- (1968): *Languages of Art: An Approach to a Theory of Symbols* (New York: Bobbs-Merrill Company).
- Grice, H. Paul (1975): “Logic and Conversation,” in P. Cole and J.L. Morgan (eds.), *Syntax and Semantics, Vol 3: Speech Acts* (New York: Academic Press).
- Happé, Francesca (1995): “Understanding Minds and Metaphors: Insights from the Study of Figurative Language in Autism,” *Metaphor and Symbolic Activity* 10:4, 275-295.
- (1993): “Communicative Competence and Theory of Mind in Autism: A Test of Relevance Theory,” *Cognition* 48, 101-119.
- Hobbes, Thomas (1651): *Leviathan or, The Matter, Form, and Power of a Commonwealth Ecclesiastical and Civil* (London: Andrew Crooke).
- Holyoak, Keith, and Paul Thagard (1995): *Mental Leaps: Analogy in Creative Thought* (Cambridge, MA: MIT Press).
- Keysar, Boaz, and Briget Bly (1995): “Intuitions Of The Transparency Of Idioms: Can One Keep A Secret By Spilling The Beans?” *Journal of Memory and Language* 34, 89-109.
- Lakoff, George (1987): *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind* (Chicago: University of Chicago Press).
- and Mark Johnson (1980): *Metaphors We Live By* (Chicago: University of Chicago Press).
- Langdon, Robyn, Martin Davies and Max Coltheart (2002): “Understanding Minds and Understanding Communicated Meanings In Schizophrenia,” *Mind and Language* 17:1/2, 127-148.
- Levin, Samuel (1977): *The Semantics of Metaphor* (Baltimore: Johns Hopkins University Press).
- Locke, John (1690): *An Essay Concerning Human Understanding* (London: T. Basset).
- Martinich, A.P. (1984): “A Theory for Metaphor,” *Journal of Literary Semantics* 13, 35-56.
- Murphy, Gregory L. (1996): “On Metaphoric Representation,” *Cognition* 60, 173-204
- Ortony, Andrew (1979): “Beyond Literal Similarity,” *Psychological Review* 86, 161-180.
- , Gerald L. Clore, and Allan Collins (1988): *The Cognitive Structure of Emotions* (Cambridge: Cambridge University Press).
- , Richard J. Vondruska, Mark A. Foss and Lawrence E. (1985): “Salience, Similes, and the Asymmetry of Similarity,” *Journal of Memory and Language* 24, 569-594.
- , Diane Schallert, Ralph Reynolds and Stephen Antos. (1978): “Interpreting Metaphors And Idioms: Some Effects Of Context On Comprehension,” *Journal of Verbal Learning and Verbal Behavior* 17, 465-477.
- Recanati, François (2004): *Literal Meaning* (Cambridge: Cambridge University Press).
- (2001): “Literal/nonliteral,” *Midwest Studies in Philosophy* 25, 264-274.
- Ricoeur, Paul (1977): *The Rule of Metaphor* (Toronto: University of Toronto Press).
- Rorty, Richard (1987): “Unfamiliar Noises I: Hesse and Davidson on Metaphor,” *Proceedings of the*

- Aristotelian Society* supp. vol. 61, 283-296.
- Rosch, Eleanor (1978): "Principles of categorization," in E. Rosch and B. Lloyd (eds.), *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum Press).
- Searle, John (1979): "Metaphor," in *Expression and Meaning: Studies in the Theory of Speech Acts* (Cambridge: Cambridge University Press).
- Shank, Roger and Robert Abelson (1977): *Scripts, Plans, Goals and Understanding* (Hillsdale, NJ: Lawrence Erlbaum Press).
- Sperber, Dan and Deirdre Wilson (1995): *Relevance: Communication and Cognition*, 2<sup>nd</sup> edition (Oxford: Blackwell).
- : "Loose Talk," *Proceedings of the Aristotelian Society* 86 (1986), 153-171.
- Stern, Josef (2000): *Metaphor in Context* (Cambridge MA: MIT Press).
- Tourangeau, Roger and Lance Rips (1991): "Interpreting and Evaluating Metaphors," *Journal of Memory and Language* 30:4, 452-472.
- Tversky, Amos (1977): "Features of Similarity," *Psychological Review* 84, 327-352.
- Weinreich, Uriel (1966): "Explorations in Semantic Theory," in *Current Trends in Linguistics, vol. 3: Theoretical Foundations*, ed. T. A. Sebeok (The Hague: Mouton Press).
- Winner, Ellen and Howard Gardner (1977): "The Comprehension of Metaphor in Brain-Damaged Patients," *Brain* 100, 719-727.
- Wolff, Phillip, and Dedre Gentner (2000): "Evidence for Role-Neutral Initial Processing of Metaphors," *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26:2, 529-541.