11-17-2009

Metaphors in Conversational Context: Toward a Connectivity Theory of Metaphor Interpretation

L. David Ritchie
Portland State University, cgrd@pdx.edu

Citation Details

Metaphors in Conversational Context:

Toward a Connectivity Theory of Metaphor Interpretation

David Ritchie

Portland State University

Author's Address

Department of Communication

Portland State University

Portland, OR 97214 USA

E-mail: cgrd@pdx.edu

Running Head: Connectivity and Metaphor

Abstract

The model proposed in this essay merges the cognitive and social elements of metaphor use and interpretation in the cognitive representation of *mutual cognitive environment* (Sperber & Wilson, 1986/1995) or *common ground* (Clark, 1996). The proposed principle of metaphor interpretation is based on the interaction of both vehicle and topic with the common ground. Because common ground is inherently problematic, it is rarely accurate to discuss "the meaning of" a metaphor, as if metaphors must have a single well-specified meaning. Each metaphor is interpreted in the particular communicative context in which it is encountered, and individual interpretations will not necessarily match unless the individuals' cognitive representations of the common ground are similar. The proposed model is consistent with several other theories of metaphor interpretation, but it explicitly acknowledges the need for interpretive processes to be context-specific and neurologically embodied (Ritchie, 2003c), and avoids the problem of circularity (Ritchie, 2003a).

Metaphors in Conversational Context:

Toward a Connectivity Theory of Metaphor Interpretation

The study of metaphors has often proceeded in an abstract, top-down way, based on the logical organization of language rather than on the way cognitive processes are embodied in the human brain and nervous system, and divorced from the social and communicative contexts in which metaphors occur. Although such top-down theories are often quite elegant from a logical point of view, the ultimate test must always be, whether the stipulated processes are compatible with the way humans actually think and communicate. Given the recent advances in understanding the neurological processes of perception and thought, it is somewhat surprising that little systematic attempt has been made, to develop theories of metaphor interpretation consistent with what we already know about the neurological functioning of the brain.

The outline of such a model is proposed by Ritchie (2004), as an alternative to Fauconnier and Turner's (1998) metaphors of conceptual "*space*" and conceptual "*blending*." In this essay I develop these ideas more systematically, beginning with Sperber and Wilson's (1986/1995) theory of relevance and Clark's (1996) model of conversation. I suggest a neural embodiment of Sperber and Wilson's concepts of *cognitive environment* and *context* and Clark's concept of *common ground* in terms of representation in working memory (Kintsch, 1998). I then propose a model of metaphor processing based on alteration of connections linking other elements present in working memory, including elements activated by the metaphor topic, to features of the metaphor vehicle that are relevant to these previously activated elements. The resulting model is consistent with many of the features of extant theories such as conceptual metaphor

(Lakoff & Johnson, 1980), as well as with what we know about the neurological

processes involved in language comprehension, and provides a firmer foundation for a

neurologically-based cognitive theory of metaphor.

Context, relevance, and common ground.

Sperber and Wilson (1986/1995) distinguish between *code* models of

communication, which assume that utterances and gestures convey particular and readily

identifiable, decodable, meanings (cf. Reddy, 1993) and *inference* models, which assume

that utterances and gestures are often ambiguous and require inference, based on the

context. Much of their discussion focuses on ostensive communicative acts (indirect

speech, gestures, etc.), but their discussion applies to most instances of everyday

language use, with the primary exception of instances in which terms have been given

precisely prescribed technical meanings, as in military, scientific, and legal language. In

their model, inference is based on a search for a context in which the communicative act

will be relevant, that is to say, will have maximum effect with minimum cognitive effort.

Cognitive environment. Sperber and Wilson (1986/1995) argue that

communication is achieved by changing the *mutual cognitive environment*. They define

cognitive environment as "the set of all facts that are manifest" to an individual (p. 39, p.

151), that is to say, all facts the individual can at the time represent mentally. A cognitive

environment is "mutual" to the extent that each person assumes that it is manifest to all

parties to the communicative event. Both parts of this definition pose potential problems,

which Sperber and Wilson go to considerable lengths to resolve. I will argue that most of

these problems arise from beginning with an *a priori*, top-down model of

communication, rather than with a neurally embodied, bottom-up model.

Throughout most of their discussion of interpretation, Sperber and Wilson (1986/1995) rely on the concept of *context*, first defined on page 15 as "The set of premises used in interpreting an utterance, … a subset of the hearer's assumptions about the world." They point out that "A context in this sense is not limited to information about the immediate physical environment or the immediately preceding utterances: expectations about the future, scientific hypotheses or religious beliefs, anecdotal memories, general cultural assumptions, beliefs about the mental state of the speaker, may all play a role in interpretation." A *context*, thus, would seem to consist of some subset of the thoughts accessible to a hearer (or a speaker), and the *cognitive environment* of a person at any given time might include several contexts, only one of which will be used in making sense of a particular communicative act.

The definition of cognitive environment in terms of "the set of all facts that are manifest" to an individual is problematic, inasmuch as it implies that facts are either manifest or not, either accessible or inaccessible. At any given time, a person is in principle capable of accessing virtually everything in her long-term memory, including facts as disparate as the multiplication table, the chronology of the American Civil War, and the approximate distance from Chicago to New York. Moreover, if she happens to be seated near a good encyclopedia or surfing the web, she has ready access to even more information. A small subset of this information is currently activated in her consciousness, some is not immediately activated but can be instantly cued, some may require a second or longer to access. Depending on the circumstances, various facts known to an individual may be more or less manifest, more or less accessible for use in making sense of a flow of conversation. The concept of cognitive effort implies that

contexts are unequally accessible, which contradicts the concept of *mutually manifest*, or at the least requires its further qualification.

Mutuality and Commonality. The concept of a *mutual* cognitive environment, those facts that are "mutually manifest" to two or more people, is if anything more difficult, given the fundamental impossibility of knowing for certain what another person is thinking at any given time. As Sperber and Wilson note (1986/1995, p. 45), evidence of what is manifest to other persons "can never be conclusive" since we have no direct evidence of what others perceive, how they perceive it, or what they think about it. We are left with assumptions about what assumptions others are making, and "even weaker assumptions about the assumptions they attribute to us." Unfortunately, in the subsequent argument, Sperber and Wilson repeatedly lose track of this fundamental indeterminacy, beginning with the last sentence of this same paragraph, in which they assert that "someone who knows an individual's cognitive environment can infer which assumptions he is actually likely to entertain" (p. 46). Since no-one ever *knows* another person's cognitive environment, inferences are always based on assumptions and guesses, a point acknowledged again in the 1995 Postscript.

We know that vocabulary, at least, is not innate, and children do not acquire vocabulary by memorizing definitions (Kintsch, 1998). We also know that humans frequently misunderstand each other or at best understand each other imperfectly, but in most circumstances approximate understanding is quite sufficient. When more precise understanding does matter, we also know that human beings employ social conventions for confirming their understanding or assuring a mutual basis for understanding ("Peter, do you see that crocodile?") From these facts about human communication, a

fundamentally solipsistic view would seem to follow, in which knowledge cannot be truly "common" or "shared," individuals can never know for certain how well they understand one another, and individuals must routinely interact on the basis of assumptions and guesses about each others' thoughts and intentions (Clark, 1996; Gibbs, 1987). Most of the time these guesses are adequate for our communicative purposes. (This point is also acknowledged by Sperber and Wilson in their 1995 postscript, and figures in their reformulation of the principle of relevance.)

Clark (1996) encounters similar difficulties in defining *common ground*. Clark defines common ground, in the first place, as all the knowledge, beliefs, and suppositions participants in a communicative act believe they share, and his definition is susceptible to many of the same semantic problems as Sperber and Wilson's (1986) definition of mutual cognitive environment. However, throughout his subsequent discussion, Clark explicitly recognizes the impossibility of knowing the extent to which two people's knowledge, beliefs, etc. match up, and much of his model of language use is devoted to explaining how interacting communicators evaluate, maintain, and when necessary improve the adequacy of their assumptions about common ground.

In Clark's (1996) version of the model, common ground is constituted by what participants *think* they share. As long as the communicative actions based on these beliefs are consistent, it is *as if* their knowledge, beliefs, etc. actually do match up. We assume common ground on a number of bases. We are reasonably confident that all human beings share certain kinds of knowledge and beliefs because of the human condition itself: We are all subject to the force of gravity, we are all susceptible to hunger, illness, etc. Beyond that, we infer common ground on the basis of shared

membership in cultural groups, co-presence in a particular environment, and a shared experience of past activities, including the current conversation thus far.

It is necessary to distinguish between general common experiences and immediate awareness. Two people who majored in English may assume they share a general knowledge of Shakespeare's important plays. However, Hamlet's relationship with Ophelia is not necessarily salient at any given time. Therefore, if one person wishes to make reference to Ophelia's madness, it may be necessary to remind the other of it before it will be part of the *effective* common ground. Similarly, if two friends are talking at a ball game, they both know there is a refreshment stand that sells popcorn and hot dogs, but that fact may not be salient to both of them. If one of them wishes to make reference to the refreshment stand, it may be necessary to remind the other of it in order to make it part of the *effective* common ground. Much of everyday communication is devoted to monitoring and maintaining common ground in this way (Clark, 1996).

Discrepancies in common ground (mutual cognitive environment) can originate in unrecognized differences in background knowledge and beliefs, or in unrecognized differences in the knowledge and beliefs of which participants are immediately aware. Even people who have much in common differ in the strength of cue needed to call a certain idea to mind, and the same person differs from one time to another in this respect. If a person has recently attended *Hamlet*, Ophelia's madness will come to mind much more readily, and related knowledge will be activated by much weaker cues as compared to another time, when she has not been recently exposed to Shakespeare. When inconsistencies creep in, participants will correct them if they seem material. If they do not seem material, and nothing is done to correct the inconsistencies, then one person

may realize that common ground does not include a certain idea while another person continues to believe that it does. In most cases, these minor slippages do not make a material difference in the outcome of communicative interactions, and are ignored.

The problem of infinite recursion. Sperber and Wilson (1986/1995) go to considerable effort to distinguish their definition of a *mutual* cognitive environment from previous concepts such as *shared knowledge*, in order to avoid the problem of infinite recursion (Peter knows Mary knows he sees the crocodile on the log but does Mary know that Peter knows she knows he sees it?) However, as Gibbs (1987) observes, they effectively sneak "shared knowledge" in through the back door, and the concept of mutual cognitive environment suffers from all the difficulties presented by the concept of "shared knowledge." Clark (1996) also goes to some length to avoid the problem of infinite recursion, and he encounters many of the same problems as Sperber and Wilson, for much the same reason.

There is no reason to think that communicators ordinarily even consider the possibility of a recursion, unless there is some reason to be concerned about it. In certain rare circumstances, when people are trying to second-guess each other's interpretations, they can and sometimes do get hung up in wondering, for example, "does she realize that I know she knows Fred is being unfaithful?" Similar emotionally complex situations may also lead people to be much more concerned than they ordinarily would about differences such as between knowledge, belief, and mere suspicion. ("She can't possibly *know* Fred is being unfaithful, but does she *suspect* it?")

A familiar example of a recursion chain involves irony. Suppose Peter says, "This tiramasu is terrible. I don't think you'll want your piece." If Mary responds by

looking with dismay at the dessert, Peter might then ask himself, "Did she realize I was speaking ironically?" Seeing his confusion, Mary might wonder, "Did he realize that I assume he intended his comment ironically?" But if Mary's response is consistent with Peter's expectation (a playful threat with her fork), there will be no need for either of them to second guess the other, and it may not even occur to either of them that the world could be in any state other than his having intended irony and her having understood that intent.

Most of our assumptions about common ground, like our assumptions about the physical world, are implicit, not explicit. When Mary steps through her door onto the street, she does not pause to wonder whether gravity will maintain her contact with the pavement or whether the pavement will support her weight. There is no reason for her to *assume* that gravitational attraction holds her firmly on the surface of the planet – she simply acts on this as a fact. Likewise, when the telephone rings at a normal volume and Peter is seated next to it, there is ordinarily no reason for Mary to wonder whether it is manifest to both her and Peter that the telephone is ringing. Only if something unexpected happens do we question these taken-for-granted aspects of the world. If the telephone rings and Peter fails to respond, *then* Mary may wonder whether he hears it. Depending on their history of telephone answering, she may also wonder whether he is aware that she assumes that he hears it. But this chain of assumptions is needed only when something about the situation prompts it, and according to the principle of cognitive economy it will ordinarily proceed only as far as is required to explain an observed anomaly.

Logically, we act on the basis of an infinite number of presuppositions, few of which are formulated as explicit assumptions. When she steps out her door, Mary acts on the presupposition that the pavement is actually solid, that the pavement is not a hallucination, that the sidewalk is not blocked by a pane of clear glass. When she hears Peter's comment about the tiramasu, Mary acts on the presupposition that she is awake and not dreaming, that they are both real human beings and not the product of an elaborate computer program, that he actually spoke and she did not hallucinate his comment, that he is speaking English and not some other language with similar sounds, that the dessert on her plate is actually food and not a clever imitation, that Peter is not involved in a plot to undermine her sanity, etc. Our brains lack the capacity to process even a tiny fraction of the presuppositions on the basis of which we live our lives, and there is no need to process them. These propositions are *mutually manifest* or part of *shared knowledge* or *common ground* only to the extent that either person is capable of affirming them if they are challenged by something in the flow of events.

The only assumptions we need to include in our model of communication are those necessitated by a realistic possibility that they might be otherwise. Only if she knows that the concrete in her sidewalk was recently repaired, and that it might possibly *not* be solid will Mary need to assume that it is solid. And only if Peter's comments or reactions disconfirm her expectations in some way (he speaks in a serious tone, but continues to eat his own tiramasu, or he speaks in a joking tone but pushes the plate aside) will Mary have occasion to question his intentions or to wonder whether or not he intended irony. How to stop an infinite recursion is a programmer's problem, and applies to the human brain only if we begin with the hardware/software distinction implicit in the

"*brain is computer*" metaphor. If we begin with an evolved brain model, there is no need to conceptualize cognition in terms of a program running on a mechanism, or to resolve the technical programming issues that arise from such a conceptualization.

   Relevance. A communicative act is relevant in a particular context only to the extent that it has some effect on that context, alters or restructures it in some way. Conversely, the degree of relevance is inversely proportional to the amount of cognitive effort required to process the assumption (to find the relevant context and to complete the effects on the context). Sperber and Wilson (1986/1995) insist that it is not necessary to assume that either the effect or the mental effort required to achieve it be quantifiable; all that is necessary is that the brain be able to compare the potential effect and required effort in one context to that of other contexts, in order to select the context in which maximum relevance (high effect for low effort) will be achieved.

   There is an unavoidable circularity in Sperber and Wilson's account of relevance, inasmuch as one can only know if the effect justifies the effort after the search for relevant contextual information is completed. In their 1995 postface, Sperber and Wilson address this circularity by defining relevance in terms of a search for *sufficient* effects, effects that will justify the effort required. They also acknowledge, first, that communication is not perfect, so that the search for relevance may not always succeed, and, second, that processing will occasionally proceed further, if the hearer is sufficiently motivated and capable (cf. Carston, 2002). Since motivation and capacity to process a message vary across time, situation, and persons, we should also incorporate a consideration of factors such as motivation, opportunity (freedom from distractions, for example), and even a disposition to look for obscure relevance (Petty & Cacioppo, 1981).

From the perspective of top-down theorizing, such a solution may seem untidy, since it leaves open the prospect of hearers not arriving at the speaker's preferred interpretation, or embellishing an interpretation far beyond the speaker's intention. However, as Sperber and Wilson acknowledge in their 1995 postface, our everyday experience of communication is characterized by just this sort of untidiness, and frequently requires us either to repair lapses in understanding or negotiate what an utterance will be taken to mean (see also Clark, 1996).

Working memory. Kintsch's (1998) concept of *working memory* suggests a basis for integrating the concepts of mutual cognitive environment and common ground with neural embodiment. According to Kintsch, *working memory* includes both short-term memory and elements in long-term memory that can be readily activated, either because of a lingering high level of activation (e.g., from recent access) or because of cues present in short-term memory. Thus, each participant's working memory will include all ideas recently activated as a result of the preceding communicative interaction (the conversational as well as the relational context), along with salient perceptions of the physical environment (the physical context). Slightly less accessible will be other ideas, memories, etc. that are closely linked to these elements. If something raises doubts about the degree to which perceptions and ideas are mutual, then working memory may also include a representation of the common ground, which ideas and perceptions are common and which may be unique to one or another of the participants.

All of these elements are "present," with different degrees or strengths of activation, and thus more or less available for linkage to new elements in the flow of conversation. The exact neural mechanisms are not known, but it is very likely that

interpretation involves co-activation of neural assemblages, which create or strengthen links between the co-activated assemblages. Elements in the flow of communication (non-verbal gestures and verbal back-channel utterances as well as fore-grounded utterances) will be processed by way of connections to activated elements in working memory (Deacon, 1997; Kintsch, 1998; Shapiro & Eigenbaum, 1997).

Cognitive effect can be thought of in terms of the degree to which processing a communicative act leads to restructuring the network of neural connections in working memory – creating new connections as well as altering the strength of previously existing connections. Thus, the search for relevance is a search for a set of neurally-embodied memories or ideas such that linkages to some aspects of the new idea introduced by a communicative act can be readily formed, and will precipitate a set of alterations in neural connections sufficient to justify the cognitive effort involved. This is consistent with the findings that a wide array of meanings and associations are briefly activated as each utterance is processed, and that those irrelevant to the current context are quickly suppressed, while the most relevant connections are strengthened (Gernsbacher, *et al.*, 2001; Glucksberg, 2001; Kintsch, 1998).

Multiple contexts. Another important feature of Clark's (1996) model is the recognition that the cognitive environment ordinarily includes several different types of activity, and that the same communicative act can affect multiple contexts. During a typical conversation, each individual maintains a representation of the topic, the flow of the conversation, the social event in which the conversation is taking place, and the nature of the relationships with other participants, as well as a representation of the common ground – what each person is currently aware of (Gibbs, 1987). A single

message can alter several of these representations, hence can be relevant in several ways at once. Moreover, even the qualities of the message, for example the fact that it includes a metaphor, irony, or some other trope, can alter or strengthen representations, including the representation of the relationship and of the nature of the current interaction (Gibbs & Gerrig, 1989).

   Working memory and common ground. In sum, at any given time, an individual involved in a communicative act has readily accessible, in the form of increased activation of neural assemblages (working memory), a representation of the topic, what has already been said by each participant, the overall purpose and progress of the communicative interchange, structural features of the discourse, such as layering (quotations, story-within-a-story, counterfactual suppositions, etc.), and relevant aspects of the relationships among various participants. Any and all of this may be brought to bear in making sense of elements within the ongoing stream of communication, and is subject to revision in response to new information.

   When I engage in a conversation with my wife about plans for the coming weekend, a relevant subset of what I know about her tastes and dispositions is active in working memory, along with a relevant subset of what I know about the range of activities we are considering. Each of these facts is linked to a network of other facts, emotions, motor control schemas, and so forth, along with cues that can activate additional information as it becomes relevant. I also have active in working memory a representation of facts that may not be activated in her working memory, which will provide a basis for building the needed common ground by introducing or reminding her

of these facts. Even in a relationship of long standing, representations of common ground are often erroneous, and must be repaired "on the fly."

Individuals can originate and interpret messages only on the basis of their understanding of the situation, since they have no direct access to the thoughts of others, and their understanding is often incomplete or erroneous. In order to understand how relevance is actually assessed and interpretation actually accomplished, we need to bear in mind the neural embodiment of common ground in the form of activated neural links.

Neural connections and interpretation of metaphor.

Consider a metaphor discussed at some length by Glucksberg, Keysar, and McGlone (1992; see also Chiappe & Kennedy, 2001; Ritchie, 2003a): *My job is a jail*. Glucksberg, Keysar, and McGlone claim that this metaphor is interpreted by establishing a super-ordinate category of "things that the metaphor vehicle *jail* typifies – situations that are unpleasant, confining, difficult to escape from, unrewarding, and so on" (p. 578). Chiappe and Kennedy protest that categorization is an unneeded step, and that this metaphor is interpreted by finding properties of the vehicle that are relevant to the topic, and transferring or mapping these properties onto the topic. In Ritchie (2003a), I argued that both proposals are circular with respect to many metaphors, since some of the properties associated with *jail*, for example, can only be applied to *job* after the process of metaphoric interpretation is already completed; in other examples, such as "that man is a *wolf* (Gentner & Bowdle, 2001), virtually none of the relevant qualities of "*wolf*" can be applied to *man* without prior metaphoric interpretation (Ritchie, 2003a; see also Carston, 2002).

Now let us consider how a connectivity model might explain the interpretation of a metaphor such as "*My job is a jail*." To begin with, let's assume that the "*jail*" metaphor is unfamiliar to the listener; I will return to the question of familiar metaphors later. The declaration, "*My job is a jail*," would ordinarily occur during a discussion of the speaker's job, in which the speaker has already complained about certain aspects of the job. It may also be spoken with a certain inflection, perhaps of exasperation. Thus, various facts about the speaker's job and her dissatisfaction with her job are already activated in working memory when the metaphor is uttered. Even when the metaphor is encountered apparently out of context, as an example in an academic discussion, the topic, *job*, will activate a set of associations, most of which (pays a salary, is a source of satisfaction) will be suppressed by the mention of the vehicle, "*jail*." In effect, the hearer will supply the context, based on whatever knowledge is available. This accounts for the observed longer time required to process figurative language when it is presented without context (Gerrig, 1989).

The vehicle, "*jail*," briefly activates links to an array of ideas, including physical descriptors (iron bars, sparse furniture bolted to the floor, locked door, concrete floor, armed guards) as well as scenarios, expectations, and schemata (a fixed sentence, punishment, strict schedule of activities, lack of privacy, physical confinement, being watched). Most of these ideas are densely connected to other ideas, as well as to social and emotional responses (social disapproval, feeling constrained and bored). Some of these associations will resonate with ideas that are already activated in working memory, either by the preceding conversational context or as a result of the reader's work in

supplying a context, and will be reinforced. Many of the associations will not resonate

with anything already in working memory, and will quickly fade or be suppressed.

The metaphor is interpreted by linking associated ideas with the existing contents

of working memory, and may alter the pre-existing links as well. The interpretation will

often depend on what has gone before. If the previous discussion has centered around the

speaker's feeling of being continually monitored and held to account for every minute of

the work-day, then the aspects of "*jail*" that have to do with being watched and lacking

autonomy will be reinforced and connected with other ideas from the preceding

discussion. If the previous discussion has centered around the speaker's need to continue

working at an emotionally exhausting job in order to repay educational loans, then

aspects of "*jail*" having to do with a prisoner's anticipated reactions to a fixed sentence,

constraint, and punishment may be reinforced. If the metaphor is encountered as an

example in an academic discussion, then whatever features of *job* are most salient to each

individual will influence the aspects of "*jail*" that are reinforced in that person's brain.

Occasionally it happens that *nothing* about the vehicle connects with anything in

working memory, in which case the search for relevance may involve activating contexts

from long-term memory, a process that requires considerably more effort, and is likely to

be undertaken only if the hearer is strongly motivated to understand the comment. If no

relevant context can be found or the hearer decides it is not worth the effort, the metaphor

may not be interpreted. For example, if the conversation occurs in the speaker's plushly-

furnished corner office after a two hour lunch, the listener may be unable to make any

sense at all of the complaint, "*My job is a jail*." In such a case, the listener may ask for

clarification, attempt to reinterpret the comment as an attempt at irony, or simply ignore

the comment about "*jail*" and continue with the conversation (Sperber & Wilson make a similar point in their 1995 "postface").

An action scenario metaphor: "*Sheep-dog*." Consider a conversation between a supervisor and the newly-appointed leader of a project team, in which the supervisor says to the team leader, "I want you to *sheep-dog* this project all the way." Many ideas, images, and actions are activated in working memory by the vehicle, *sheep dog*. These may include the image of a dog running about behind a flock of sheep, the dog's barking, the sounds and smells of the sheep, and perhaps a green meadow. More specific images may also be activated, such as sheep bolting from the flock and the dog nipping at their legs to chase them back into the group and keep them moving, or images of a predator such as a fox or coyote stalking the flock, and the dog chasing it away.

None of the direct associations activated by the phrase *sheep dog* will resonate with anything already activated in working memory; these immediate associations will have little effect on the activation levels of working memory. Some of the associations with *sheep dog*, however, will include links to expectations, emotions, and scenarios that will resonate with contents of working memory. For example, secondary associations with a *flock* of sheep, including the apparently random behavior of individual sheep, may resonate with secondary associations with a *team* of assorted engineers, accountants, and so forth; connections to these attributes are likely to be reinforced, as they are linked to the resonant attributes associated with *team*. The sheep dog's herding activities, which are directly linked to the random behavior of the sheep, will also be reinforced in working memory, in turn facilitating indirect connections to activities of a team leader with respect to team members. In the "*sheep dog*" metaphor, some of the indirect, secondary

associations of the vehicle link up with indirect, secondary associations of the topic; once these indirect links are made, the initial interpretation of the metaphor is established, and the links between *team leader* and *team* are altered to fit this new information.

Contextualization and depth of processing. Depending on the nature of the conversation, the processing of the metaphor may stop with the initial, relatively simple, interpretation ("He wants me to keep the group working toward common goals, just as a sheep dog keeps a flock moving in a common direction.") However, if the team leader's relationship with her supervisor or other aspects of the conversation context warrant, further processing may lead to elaboration of metaphorical entailments. For example, other qualities associated with a flock of sheep include the presence of coyotes, foxes, or other predators, so *protection from predators* may also be linked to *sheep dog*. Similarly, in an organizational context in which teams compete for resources, the concept of a project team may link to the concept of competition among work units for resources. Secondary associations (expectations, emotions, etc.) with both sets of concepts (*protect from attacks* and *preserve team resources*) will stimulate links between these attributes of topic and vehicle, and generate a metaphorical entailment in which part of the role of team leader is "*protecting*" the team against "*attacks*" from other managers. Further processing may lead to links connecting the team leader's relationship to her supervisor with the sheep dog's relationship to the shepherd, who trains and controls the sheep dog. Thus, the choice of the "*sheep-dog*" metaphor may entail limiting the team leader's autonomy and retaining overall control of the project in the hands of the supervisor.

How many of these entailments will be processed will depend in large part on the listener's representation, in working memory, of the communicative situation, her

relationship with her boss, and so on. In Sperber and Wilson's (1986/1995) model, there

is some optimal degree of processing, in which the maximum amount of cognitive effects

are achieved for a given amount of cognitive effort. However, Petty and Cacioppo's

(1981) "elaboration likelihood model" suggests that the degree to which a message is

processed is influenced by ability factors, including available time, lack of distractions,

and background knowledge as well as by motivation factors, including perceived

importance and what they call "need for cognition" – an intrinsic enjoyment of thinking

for its own sake. (For example, a cognitive linguist will often continue processing a

metaphor long after everyone else has forgotten it.) If the hearer has no associations

whatsoever with the phrase *sheep dog*, then ability to process will be minimal, and the

hearer may either create an idiosyncratic interpretation through some imaginative

connection, or find it impossible to make any sense of the metaphor, and treat it as

uninterpretable jargon.

Connecting topic and vehicle: The role of conversational context.

Consider a domestic metaphor: "My wife is an *anchor*." The word *anchor* is

connected with a dense array of perceptual images, motor control schemata, and abstract

concepts. The perceptual associations include weight, iron rust, shape, boats and ships,

open water, bottom muck, and many others. Slightly more abstract associations include a

cessation of movement, constraints on movement to within a small circle with a radius

defined by the length of anchor chain, safety from grounding on a shoal, security from

undesired movement, relaxed vigilance, remaining in a harbor, etc. Associations may

also be activated that link *anchor* to remembered boating scenes, scenes from fiction or

movies, an uncle's stories about his war-time service in the navy. None of these has any

connection with *wife*. So how might a listener make sense of such a metaphor?

Recall that each individual maintains, in working memory, a mental

representation of the conversational context, a web of activated connections among the

pertinent elements of the conversation. Activated connections in working memory

include background knowledge and intentions as well as most of what has already

transpired in the present conversation. In order to arrive at an account of how a listener

might make sense of the statement that "*My wife is an anchor*," we need to consider a

conversation in which such a statement might occur (Gibbs & Gerrig, 1989).

Conversational Context. Imagine a conversation between two friends, after a

long separation, in which they bring each other up to date on their experiences since they

last saw each other. Given these contextual factors, it is reasonable for each participant

in the conversation to assume that both participants have active in working memory a

similar set of connections to existing knowledge about the friendship and the relational

and emotional history of each person, along with whatever information has been

introduced during the conversation thus far. All of these are continually revised, up-

dated, and connected to new information as the conversation progresses, and some but

not all of these new connections will become part of permanent memory.

Now imagine two versions of the conversation:

(1)   "You seem much happier than the last time I saw you. You used to be

discontented and easily distracted, but now you seem to be contented and at peace

with yourself."

"My wife is an *anchor*."

Or, an alternative conversational context,

> (2)   "You sound like you've become bored with life. You used to be so eager
>
> for new experiences, but now the old zest for life seems to have become
>
> dulled."
>
> "My wife is an *anchor*."

Since none of the qualities directly associated with *anchor* has any apparent connection

to *wife*, the connection must be made on the basis of secondary linkages. On a ship or

boat, being *at anchor* is associated with a state of relaxed vigilance and a feeling of

safety, an opportunity to catch up on deferred maintenance and make plans for the next

leg of a voyage, and indirectly with a sense of restfulness and general well-being.

Sometimes being *at anchor* is also associated with lack of progress and impatience to be

underway, and indirectly with boredom and frustration.

In the context of conversation (1), working memory includes a contrast between a

previous state of discontentment and distraction, and a current state of contentment and

peace. Ideas and emotions associated with security, relaxed vigilance, and safety will

connect with the ideas and emotions associated with contentment and lack of worries in

the speaker's current life, already activated in the common ground, so will be

strengthened and connected to the concept of *wife*, thereby creating or strengthening

connections between *wife* and feelings of contentment and lack of worries. Ideas and

emotions associated with frustration and a lack of progress toward a goal will be

suppressed because they are inconsistent with the ideas already activated in working

memory. In the context of conversation (2), working memory includes a contrast

between a previous zest for life and a current state of boredom, so the pattern of

connections will be just the opposite as in the first conversation. In both cases, the ideas and emotions activated during this interpretive process will be connected to the similar ideas and emotions previously activated in the participants' working memories, and will remain as part of each participant's working memory, where it may influence processing of subsequent information (Allbritton, McKoon, & Gerrig, 1995).

To recap the interpretive process, each utterance is processed in the context of the currently activated information, and each participant's representations are updated accordingly, including representations of the conversation thus far and of the underlying topic and purpose. When the speaker mentions "my wife," the activation level of features associated with the concept of wife is increased, along with the activation level of any other information about the speaker's relationship to his wife that may have been introduced into the conversation. When the speaker mentions "*anchor*," diverse elements (facts, ideas, images, and emotions) associated with "*anchor*" are momentarily activated; those that afford no ready connection with the contents of working memory are suppressed, and those that resonate with the contents of working memory are reinforced. Qualities such as shape, color, weight, and metallic composition of an anchor connect with nothing in either conversation, and will be quickly suppressed.

Metaphorical entailments. Metaphorical entailments may become activated, as neural links to related concepts, and enter into the network of meanings that is constructed and maintained during a conversation in at least two ways. In the instance of the "*anchor*" metaphor, for most adult speakers of English, there already exists a set of metaphors in which emotional commitment is expressed as "*attachment*," a lack of purpose or commitment as "*drifting*," a disruption of plans as "*shipwreck*" or "*running*

aground," and so on. These stock metaphors may be part of the general cultural

knowledge of both participants, exist in both brains in the form of neural links among the

related concepts, and hence may be assumed to be accessible as part of the extended

common ground (cf. Lakoff & Johnson, 1980). But if they are not already known, part of

the accessible common ground, the metaphorical connections may need to be generated

through extended processing in which secondary neural links between attributes of topic

and vehicle are established or strengthened (Gerrig, 1989; Gerrig & Murphy, 1992).

       The elaboration of the "*anchor*" metaphor into extended entailments can be

visualized as follows. First, the network of neural nodes associated with *wife* and the role

of a wife with respect to a man's life is connected with the network of neural nodes

associated with *anchor* and the function of an anchor with respect to a ship, by way of

links to emotions associated with concepts such as safety and security. (In other cases the

links might involve perceptions or motor activities; the principle is much the same.)

Then, if further processing seems warranted, this single link or small set of links may

expand outward, as other connections are discovered: For example, *anchor* is negatively

associated with *running aground*, which is associated with emotions such as disruption of

plans, frustration, and a sense of failure, which are then negatively associated with *wife*,

who is an "*anchor*," in the sense that she prevents the speaker from "*running aground*."

       As a metaphor is interpreted, new neural connections are formed between the

network of neural connections associated with the vehicle (*anchor*, as well as related

concepts such as *ship* and *harbor*) and the neural connections associated with the topic

(*wife*). Of at least equal importance, the metaphor can also alter the connections between

the topic and the other elements already present in the cognitive context, in this example,

by strengthening the link between *wife* and the speaker's *quality of life*. No longer is the

speaker's relationship with his wife a merely coincidental part of his overall sense of

well-being (conversation 1) or ennui (conversation 2); the "*anchor*" metaphor moves her

into a central and overall causal role in the picture the speaker is painting of his current

life. The metaphor strengthens the connections between the speaker's wife and other

facts that have already been introduced into the conversation, and at the same time lays

the foundation for connecting her to aspects of his life that have yet to be mentioned.

(Compare "*My wife is an anchor"* to "*My career is an anchor*.") The metaphor may also

connect the person concept of the speaker to ideas associated with the relationship of ship

to anchor, such as "*safe from drifting.*" As a secondary effect, these changes may also

precipitate changes in the hearer's understanding of the conversation itself, for example

changing it from casual chit-chat to a more intimate sharing of emotional confidences

(Gibbs & Gerrig, 1989).

    The extent to which metaphoric entailments are elaborated is partly determined,

as noted in the foregoing, by elements already activated in working memory, such as the

purpose of the conversation and the degree to which the metaphor is emphasized by the

speaker. The search for optimal relevance, as posited by Sperber and Wilson

(1986/1995), plays a role, but by itself it cannot fully explain differences between

individuals and within individuals between situations in the degree to which metaphoric

interpretations are elaborated. The elaboration likelihood model (Petty & Cacioppo,

1981) fills in some of the gaps, by suggesting factors related to both ability and

motivation that influence the degree to which a message is processed. If the pace and

timing of the conversation allows reflection on the metaphor, more entailments are likely

to be considered, and the connections are likely to be strengthened even further through rehearsal. But the hearer's understanding of the communicative situation itself can also constrain elaboration (Clark, 1996).

Constraints on elaboration. Should unauthorized entailments occur to the listener, the listener's perception of the common ground will shape the way these are processed. The constraints on elaboration of a metaphor such as "*anchor*" or "*sheep-dog*" may be social as well as cognitive. Within the conversation between a team leader and supervisor, elaborating the metaphor along lines that do not advance the apparent, shared, purpose of the conversation (for example, by commenting that the team members are indeed "a *wooly*" bunch) would usually be poorly received.

Should some of the entailments seem relevant to other ideas the listener wishes to introduce into the conversation, the common ground will need to be modified in the listener's next speaking turn, so as to render the entailments salient to the other person as well: "Well, just be sure you don't get too comfortable in this *safe little harbor*." Such an attempted expansion of metaphorical entailments may lead to negotiation over what the metaphor is to be taken to mean (Clark, 1996). Excursions into metaphorical entailments that are totally inconsistent with the intent of the speaker ("You know, anchors work by getting themselves *buried in the muck*") might be resisted strongly enough to disrupt the conversation entirely.

Ambiguous Metaphors. The ideas developed in the foregoing are similar to Lakoff and Johnson's (1980; 1999) claim that the topic is actually *experienced* in terms of the vehicle, but the connectivity account places more emphasis on the contextual relativity of the metaphor, and does not assume that a conceptual metaphor precedes the

linguistic expression. Moreover, because the connectivity model starts with the act of interpreting a particular message, it does not share the assumption of conceptual metaphor theory that thematically similar expressions are necessarily expressions of a common underlying conceptual metaphor. Nor does the connectivity model lead to the hypothesis that a speaker or writer will be more likely to draw different expressions for a single topic from a common conceptual metaphor, as does conceptual metaphor theory (Shen & Balaban, 1999; Keysar & Bly, 1999).

In the "*anchor*" examples, the hearer is invited to understand the role of the wife in the speaker's overall life in terms of the role of an anchor with respect to a sailor. In each of the conversations, the metaphor moves the wife closer to the center of the speaker's account of his current life, but in each conversation, "*anchor*" and *wife* are connected to distinct sets of ideas, so the metaphor restructures the cognitive context in a very different way in each example. In the first conversation, the metaphor credits the wife with the speaker's satisfying and productive life. In the second conversation, the metaphor blames her for the speaker's restlessness, boredom, and ennui. The same metaphor plus different conversational and cognitive contexts leads to different neural connections, hence to different meanings, not only for the metaphor but also for the metaphorically described concept, *wife*.

Prior Knowledge of the Topic

Vervaeke and Kennedy (in press) claim that we can only interpret and evaluate a metaphor on the basis of prior knowledge about the target or topic, and that indeed a metaphor can only exist if the topic has a pre-existing structure. But the connectivity model would suggest that a metaphor can be understood with no prior knowledge of the

topic, provided that there is sufficient information in the context to guide interpretation,

and that the interaction of metaphor and context can provide a good initial understanding

of a previously unknown topic. For example, until I visited Australia a few years ago, I

had never heard of *barramundi*. I don't remember how I first encountered the dish, but it

is easy to imagine the following conversation in a Sydney restaurant:

> Waiter: "Have you ever had barramundi?"
>
> Ritchie: "What is barramundi?"
>
> Waiter: "It is the *beef cow* of the northern rivers and estuaries."

The context, a conversation with a friendly waiter about what to order for dinner,

provides the information that we are talking about a food item; the phrase "of the

northern rivers and estuaries," in the vehicle, tells me that we are not talking about cattle,

but about something that lives in tropical waters. From this information it is possible to

infer that the topic of the metaphor is some kind of aquatic creature, and to eliminate

many properties associated with *beef cow* without further consideration (e.g., has horns,

hooves, and hair; says "moo"). Properties of *beef cow* that seem relevant in the context

include "is hearty and flavorful," "is considered good to eat," and perhaps "is relatively

large." Even with no prior knowledge of the topic, the context provides the information

needed to arrive at a reasonable interpretation (Ritchie, *in press*). In neurological terms,

the context plus the phrase, "of the northern rivers and estuaries," suppresses some links,

generates or strengthens others, and alters the customer's cognitive context to become

more consistent with that of the waiter.

Even in an example such as this, the baseline assumption of common ground can

fail, leading to a breakdown in communication. Suppose the waiter is an ecologically-

minded chap who believes that the barramundi fishery is destroying the ecology of the

coastal estuaries. "Barramundi is the *beef cow* of coastal estuaries" might then be

intended to refer to the similar environmental effects of cattle ranching and barramundi

fishing, rather than to the culinary qualities of these animals, and the waiter's intent might

be to encourage the customer to join in a boycott of barramundi, perhaps by ordering a

vegetarian dish. Here it is useful to point out that the *actual* effects of cattle ranching or

barramundi fishing are irrelevant; what matters are the waiter's and customer's cognitive

representations of their own and each others' beliefs (Ritchie, *in press*).

Why use metaphors?

The foregoing account suggests at least a partial answer to the question, "Why do

people use metaphors?" Both "*sheep-dog*," used as a verb, and "*anchor*," in a particular

conversational context, activate dense webs of connections to other ideas, expectations,

responses, and scenarios, ready to connect to previously activated elements in listeners'

working memory. Metaphors are not always deeply processed, and sometimes they serve

as little more than stock expressions (Gentner & Bowdle, 2001). Thus, the team leader in

the "*sheep-dog*" example may not process the metaphor beyond a simple "keep the group

working toward common goals." But if the conversational context and tempo allows

further processing, it may catalyze an extensive re-structuring of the way she understands

her newly assigned role.

In conversation (1) of the "*anchor*" example, the speaker might have used any of

a number of other metaphors to express the idea that his wife has been a central causal

agent in his current state of contentment, such as "*beacon*," "*cheer-leader*," or "*spark-

plug*." Each of these alternatives carries very different entailments, and, if processed

deeply, each would restructure working memory in a distinct way. The "*anchor*"

metaphor casts the wife in the role of a stabilizing influence and the speaker in the role of

a volatile, mercurial sort – "*My wife is a spark-plug*" would have the opposite

entailments. Thus, a well-chosen metaphor can potentially convey a rich array of ideas

and impressions, including quite subtle nuances of thought and feeling. (As Gibbs and

Gerrig (1989) point out, the mere fact of using a metaphor can in particular create or

render salient a heightened level of intimacy). In most cases, other than the most trite or

hackneyed metaphors, use of an apparently simple metaphor can serve to precipitate

extensive changes to the listener's cognitive context. Religious metaphors such as "*God

the father*" and "*rebirth*" can precipitate a restructuring of cognitive connections that

reaches far beyond those activated in the hearer's working memory at any one time: That

is why they are so culturally, intellectually, and emotionally powerful.

Connections to other models.

        The primary features that distinguish the account I have just described from other

accounts of metaphor are the heavy reliance on assumed common cognitive activation of

knowledge, images, emotions, etc. for selection and interpretation of metaphor and the

implication that interpretation is often indeterminate and potentially problematic. The

connectivity model emphasizes, more than most theories of metaphor, the importance of

conversational context, broadly defined to include all information activated in working

memory. The concepts of relevance (Sperber & Wilson, 1986/1995) and common

ground (Clark, 1996; Gerrig & Murphy, 1992; Gibbs, 1987; Gibbs & Gerrig, 1989) unite

social and cultural processes of communication with cognitive processes, inasmuch as

each individual has a cognitive representation of common ground and bases interpretation

on that representation of common ground, and inasmuch as maintaining the common

ground is a continuing task throughout a conversation.

The connectivity model is consistent with theories such as categorization (Keysar

& Glucksberg, 1992), trait mapping (Gentner & Bowdle, 2001), and trait attribution

(Chiappe & Kennedy, 2001); for a detailed discussion, see Ritchie (2003a; 2003b).

However, the connectivity model avoids the problem of circularity inherent in each of

these theories (Ritchie, 2003a) by specifying a mechanism in the form of neural

connections in active working memory by which the relevant attributes of the vehicle

(those consistent with elements of active working memory) are selected and connected to

relevant attributes of the topic.

Conceptual Blending. The connectivity model also has some resemblance to

Conceptual Integration or Conceptual Blending (Fauconnier & Turner, 1998; 2002),

inasmuch as the connectivity model posits a connection at the neural level between

context-relevant aspects of the separate concepts associated with topic and vehicle.

Unlike Conceptual Blending Theory, the connectivity model does not posit duplication of

conceptual elements in a separate cognitive "*space,*" and the connectivity model places a

much stronger emphasis on the cognitive representation of common ground, already

activated in working memory when a metaphor is encountered. Conceptual Blending

Theory posits a need for common elements in two concepts as a basis for integration, and

thus cannot explain an example such as "*My wife is an anchor*" in which topic and

vehicle have no qualities whatsoever in common (Ritchie, 2004). The connectivity

model suggests that any two concepts can be integrated on the basis of common links to

elements including attributes, emotions, and expectations that are activated in working

memory. Thus, the connectivity model is capable of explaining abstract metaphors such as "*My wife is an anchor,*" as well as the tropes cited by Fauconnier and Turner.

Because conceptual blending posits the creation of a separate cognitive space for each blend, it implies that processing a blend would take longer than processing a simple statement. Moreover, processing a complex blend, in which two blended concepts are in turn blended, should take much longer than processing a non-figurative statement of similar length and complexity. For example, "What this country needs is *Margaret Thatcher* for President" should, according to conceptual blending theory, take twice as long to process as "What this country needs is *Oliver North* for President." Because it posits only alterations in relevant connections among the neural representations of various concepts, the connectivity model does not imply greater processing time for messages that require integration of separate concepts than for messages that do not require conceptual integration. For a more detailed discussion, see Ritchie (2004).

Further Implications. A major advantage of a connectivity model is that it lays the foundation for continued incorporation of new tools and concepts from neurological research into our understanding of metaphor. Recent advances in techniques for mapping brain activity during the processing of sensory stimuli might effectively be used to test theoretical hypotheses about, for example, the differences between novel, familiar, and fully lexicalized metaphors. Processing of language frequently involves activation of multiple regions of the brain, including areas specializing in visual, auditory, and tactile perceptions (Deacon, 1997; Jung-Beeman *et al*., 2004; Marcus, 2004). When the language is metaphorical, areas in the right hemisphere associated with abrupt insights (the "aha!" effect) in problem solving (Bottini, *et al*., 1994; Jung-Beeman *et al*., 2004)

are also activated. If a metaphor such as "Her theory *lighted up* the dark corners of our

understanding" has become fully lexicalized, we would expect little activation in the

sensory processing areas of the brain while it is processed, but if the hearer actually

experiences the topic in terms of the vehicle, we would expect a higher level of activation

in brain areas committed to processing visual stimuli.

Concluding Remarks

The model proposed in this essay draws from several theories of language use,

communication, and metaphor, notably including Conceptual Metaphor Theory (Lakoff

& Johnson, 1980), Sperber and Wilson's (1986) theory of relevance, Clark's (1996)

theory of common ground in conversation, and Kintsch's (1998) connectionist theory of

language, as well as his concept of working memory. In the connectivity model the

cognitive and social elements of metaphor use and interpretation are merged in the

cognitive representation of common ground. The primary principle of metaphor

interpretation in the connectivity model is to be found, not exclusively in the vehicle or

topic of a metaphor, but in the interaction of vehicle and topic with the common ground.

Common ground is inherently problematic, and the establishment and

maintenance of common ground is at least implicitly a function of all communicative

acts. An implication is that interpretation and understanding are also inherently

problematic. Specifically with respect to metaphors, an implication of the connectivity

model is that it is rarely accurate to speak or write of "the meaning of a metaphor," as if

metaphors could have a single well-specified meaning. Rather, a particular metaphor in a

particular communicative situation has an intended (speaker's) interpretation and is

independently interpreted by each listener: If the participants' assumptions about

common ground are valid, and each has a similar set of ideas activated in working memory, then their interpretations are likely to be quite similar. If participants' cognitive contexts do *not* match, their interpretations may be very different from one another.

A methodological implication of the view proposed herein is that metaphors, as communicative elements, cannot reasonably be discussed in the absence of their communicative contexts. When theorists do discuss metaphors in an a-contextual way, an implicit context must be supplied by all participants in the discussion. Unless every participant can be counted on to supply a reasonably similar context, there is no guarantee that they will come to similar conclusions about the meaning of the metaphor. For a similar point, see Keysar and Bly's (1999) recent work on ambiguous metaphors.

The connectivity model also recognizes that metaphors are often ambiguous and sometimes simply aren't interpreted at all, and suggests an account of why this might be so. Most importantly, the connectivity model is explicitly grounded in a language-based theory of communication that is consistent with much of what we know about communication, both at the level of neurological processing and at the level of social interaction and cultural conventions.

References

Allbritton, D. W; McKoon, G.; Gerrig, R. J. (1995). Metaphor-Based Schemas and Text Representations: Making Connections through Conceptual Metaphors *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 3, 612-625.

Bottini, G., Corcoran, R., Sterzi, R., Paulesu, E., Schenone, P., Scarpa, P., Frackowiak, R. S. J., and Frith, C. D. (1994). The role of the right hemisphere in the interpretation of figurative aspects of language: A positron emission tomography activation study. *Brain, 117*, 1241-1253.

Carston, R. (2002). *Thoughts and utterances: The pragmatics of explicit communication*. Oxford, England: Blackwell.

Chiappe, D. L., and Kennedy, J. M. (2001). Literal bases for metaphor and simile. *Metaphor and Symbol, 16*, 249-276.

Clark, H. H. (1996). *Using language*. Cambridge, England: Cambridge University Press.

Deacon, T. W. (1997). *The symbolic species: The co-evolution of language and the brain*. NYC, NY: W.W. Norton.

Fauconnier, G., and Turner, M. (1998). Conceptual integration networks. *Cognitive Science, 22*(2), 133-187.

Fauconnier, G., and Turner, M. (2002). *The way we think: Conceptual blending and the mind's hidden complexities*. New York, NY: Basic Books.

Gentner, D., and Bowdle, B. F. (2001). Convention, form, and figurative language processing *Metaphor and Symbol, 16*, 223-247.

Gernsbacher, M. A., Keysar, B., Robertson, R. W., and Werner, N. K. (2001). The role of suppression and enhancement in understanding metaphors. *Journal of Memory and Language, 45*, 433-450.

Gerrig., R. J. (1989). The time course of sense creation. *Memory & Cognition, 17*, 194-207.

Gerrig, R. J, and Murphy, G. L. (1992). Contextual Influences on the Comprehension of Complex Concepts. *Language and Cognitive Processes, 7*, 3-4, 205-230.

Gibbs, R. W., Jr. (1987). Mutual knowledge and the psychology of conversational inference. *Journal of Pragmatics, 11*, 561-588.

Gibbs, R. W., Jr., and Gerrig, R. J. (1989). How Context Makes Metaphor Comprehension Seem "Special." *Metaphor and Symbolic Activity, 4*, 3, 145-158.

Glucksberg, S. (2001). *Understanding figurative language: From metaphors to idioms. Oxford Psychology Series Number 36*. Oxford, England: Oxford University Press.

Glucksberg, S., Keysar, B., and McGlone, M. S. (1992).Metaphor understanding and accessing conceptual schema:Reply to Gibbs(1992).*Psychological Review, 99*, 578-581.

Jung-Beeman, M, Bowden, E. M., Haberman , J., Frymiare, J. L., Arambel-Liu, S., Greenblatt, R., Reber, P. J., Kounios, J. (2004). Neural Activity When People Solve Verbal Problems with Insight. *PLOS Biology 2* (4*).

Keysar, B., and Bly, B. M. (1999). Swimming against the current: Do idioms reflect conceptual structure? *Journal of Pragmatics, 31*, 1559-1578.

Keysar, B., and Glucksberg, S. (1992). Metaphor and communication. *Poetics Today, 13*(4), 633-658.

Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. Cambridge, England: Cambridge University Press.

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

Lakoff, G. and Johnson, M. (1999). *Philosophy in the flesh: The embodied mind and its challenge to Western thought*. New York: Basis Books.

Marcus, G. (2004). *The birth of the mind: How a tiny number of genes creates the complexities of human thought*. New York, NY: Basic Books.

Petty, R. E., and Cacioppo, J. T. (1981). *Attitudes and Persuasion--Classic and Contemporary Approaches*. New York, NY: McGraw-Hill.

Reddy, M. J. (1993). The conduit metaphor: A case of frame conflict in our language about language. In Ortony, A. (ed.), *Metaphor and Thought, 2ⁿᵈ Ed*, pp. 164-201. Cambridge, England: Cambridge University Press.

Ritchie, D. (2003a). Categories and similarities: A note on circularity. *Metaphor and Symbol, 18* , 49-53.

Ritchie, D. (2003b). "ARGUMENT IS WAR" – Or is it a game of chess? Multiple meanings in the analysis of implicit metaphors. *Metaphor and Symbol, 18*, 125-146.

Ritchie, D. (2004) Lost in space: metaphors in conceptual integration theory. *Metaphor and Symbol, 19*, 31-50.

Ritchie, D. (*in press*). Common Ground in Metaphor Theory: Continuing the Conversation. *Metaphor and Symbol*.

Shapiro, M. L., and Eigenbaum, H. (1997). Learning and memory: Computational principles and neural mechanisms. In Rugg, M. D., Ed. *Cognitive Neuroscience*, pp. 77-130. Cambridge, MA: MIT.

Shen, Y, and Balaban, N. (1999). Metaphorical (in)coherence in discourse. *Discourse Processes, 28*, 139-153.

Sperber, D., and Wilson, D. (1986/1995). *Relevance: Communication and cognition*. Cambridge, MA: Harvard University Press.

Vervaeke, J. and Kennedy, J. M. (in press). Conceptual Metaphor And Abstract Thought. *Metaphor and Symbolic Activity*.