From *The Literal and Nonliteral in Language and Thought.* Edited by Seana Coulson and Barbara Lewandowska-Tomaszczyk. Frankfurt: Peter Lang, 2005. Pages 25-52.

This chapter is an abbreviated excerpt from a chapter in Cristina Cacciari, Ray Gibbs, Jr., Albert Katz, and Mark Turner. *Figurative Language and Thought.* New York: Oxford University Press, 1998

The Literal Versus Figurative Dichotomy

In previous work (Turner 1989, 1991, and 1996b), I offered demonstrations that the commonsense dichotomy between "literal" and "figurative" is a psychological illusion. There is no doubt that some products of thought and language feel literal while others feel figurative. We have reactions, and they are motivated, but these motivations do not come from fundamental differences of cognitive operations. "Literal" and "figurative" are labels that serve as efficient short-hand announcements of our integrated reactions to the products of thought and language; they do not refer to fundamentally different cognitive operations.

The commonsense dichotomy between "literal" and "figurative" arises from a folk theory concerning thought, reality, and language, or more technically, entities, categories, reference, predication, truth-conditionality, and compositionality: in this folk theory, an entity (this tree outside my window) is a bundle of features (photosynthesis, etc.); a category (tree) is a bundle of criterial features (trunk, limbs, roots, photosynthesis, etc.) shared by its members; common nouns ("tree") refer to categories of objects; verbs ("grow") refer to categories of events; verbs predicate event-features of their subjects ("a tree grows" predicates *grow* of *tree)*; adjectives ("big") and adverbs ("slowly") modifying common nouns and verbs assign or remove features (big trees grow slowly); and predication and assignment are compositional in the sense that the meaning and truth-value of any conjunction is just the conjunction of the meanings and truth-values of the components, so that assigning a complex feature is no different from assigning the set of its component features. For example, the subject of "big trees grow slowly" refers to the subcategory of *tree* whose members are additionally *big* (i.e., all objects of which it is true both that "this is big" and "this is a tree"); the verb phrase refers to a subcategory event that has the feature *slowly* along with all the features of the event-category *grow;* and the entire sentence predicates the features of *grow slowly* of the subcategory *big tree*; that is, it adds the features of *grow slowly* to the features of *big tree*. This addition is compositional for both truth-conditions and meaning.

In this folk theory, a connection is true if the state of affairs to which it refers is the case in the world (i.e. it is the case that everything that is both *big* and a *tree* has the complex feature *grows slowly*). It is false if the state of affairs to which it refers is not the case in the real world.

In this view, *The sun is a useful star* predicates of the sun both the feature *useful* and all the features of the category *star.* This predication composes the

features of *the sun*, *useful*, and *star*. The assignment of the complex feature *useful star* is just the composition of the assignments of its component features; all of the component assignments are true (there is a sun with all of its features and it has additionally the feature *useful* and all of the features of *star*), and so their composition is true; the connection is therefore true; the thought is therefore true; and a statement of that thought is therefore true.

"The sun is a planet," in this view, means that it is the case that there is a sun with all its features and that it has all the features of the category *planet*. The thought and expression are therefore viewed as false. (I pass over more sophisticated machinery needed to talk about cases like "If I were English, I would drink tea instead of coffee.")

In this commonsense folk theory, thought and language operate by truth-conditional composition of features. This is the realm of "literal" thought and language. It is a consequence of this view that there must be a separate kind of thought and language, called "figurative," that uses alternative cognitive operations. The logic that leads to this consequence runs as follows: "The sun is a jewel" is just literally false; its literal meaning is a composition of the features of *the sun* and *jewel*, and that composition is not the case in the world. This composition exhausts all the literal meaning and all the literal truth-value the expression can have. Therefore, if "The sun is a jewel" has any alternative meaning or any alternative truth-value, it can have it only by virtue of some different process, of interpretation or of conceptual connection. Since everyone recognizes that "The sun is a jewel" can mean something aside from "the sun exists and it has all of its own features as well as all the criterial features of the category *jewel*," and since everyone recognizes that "The sun is a jewel" can moreover have (at least something like) positive truth-value, we must conclude that there is some alternative process, some "figurative" process, by which it acquires this alternative meaning and truth-value. "Figurative" here means exactly "not literal."

Clearly, we have different reactions to "The sun is a star" and "The sun is a jewel." We feel that an important difference between them is signified when we call the first "literal" and the second "figurative." No question. What is at issue is whether these different reactions indicate fundamentally different cognitive operations in the different cases. I have proposed that the answer is No.

In my previous work, I have proposed that conceptual connections between two mental arrays strike us differently depending on how those arrays are already related in our category structures. A connection seems literal or figurative (or somewhere in between) not absolutely but in relation to the category structures used to understand it. "A child is a light bulb" asks us to connect mental arrays that are basic level categories, and thus seems figurative. "Parsley is cumin" or "A mug is a glass" or "A steno chair is a rocking chair" asks us to make the same kinds of connections between mental arrays, but in these cases the two mental arrays share a supercategory at or below the basic level (e.g., chair), and so we feel that they are literal (but false). In cases like "Parsley is cumin," the usual claim does not hold that recognizing literal falsity prompts us to recognize "figure" as a way of repairing the falsity. In all these cases, the

feeling that something is literal or figurative depends not on special mechanisms of connection but rather upon the relative status of the elements connected.

I have also argued that there is another, related influence on judging a connection to be literal or figurative: the degree to which the conceptual connection or the linguistic expression is generatively entrenched. The greatest degree of generative entrenchment for a conceptual connection occurs when it becomes established as a central part of basic category structure: for example, a woman is a human being. But there are other conceptual connections between elements in category structures that, while not sufficiently generatively entrenched to seem to belong to our "literal" categories, are nonetheless available to us—a woman is a vessel, for example, Generative entrenchment of mental connection is a graded scale. We connect *wind* to *intentional agent*, *life* to *drama*, and an object's stasis on a table with the action of *holding* something *up*, all with varying degrees of generative entrenchment. "Life is metabolism," "Life is a performance," "Life is a play," "Life is a cast of dice," and "Life is an isosceles triangle" all ask us to locate conceptual connections that differ in their degree of generative entrenchment in our conceptual systems. Our reactions to these expressions differ accordingly. "Life is metabolism" sounds (to me) literal and definitional; "Life is a play" sounds (to me) half-way between literal and figurative; "Life is a cast of dice" sounds figurative and commonplace; "Life is an isosceles triangle" sounds wildly figurative. It also sounds unintelligible to me, until I finally find a connection: Life is like an isosceles triangle; it always has its irregular side.

Consider "I am making intellectual progress." This expression depends on the conceptual connection between *a thinker* and *a person moving in space*, analyzed by Eve Sweetser (1990). When we think about it, this connection does not seem to us to belong to our "literal" category structures—a thinker is not "literally" a traveler. Yet the connection is so entrenched as to be immediately and automatically available from the conceptual domain of *thinking*: no conceptual work is needed to build the connection; the connection to *moving in space* does not need to be activated for new inferential or semantic work. Moreover, connections of this sort typically bring along entrenched grammar and vocabulary: "intellectual progress" follows a standard grammatical pattern for connections in which the adjective comes from the domain to which we wish to refer (*thinking*) while the noun comes from the other domain (*moving in space*). The lexical filling ("intellectual progress") of this grammatical pattern is also highly entrenched. Accordingly, the connection and the expression can strike us as literal.

"Mental journey" strikes us slightly differently. "Intellectual progress" and "mental journey" depend on the identical conceptual connections expressed in the identical phrasal pattern (adjective from the domain referred to, noun from the other domain), but the vocabulary of "mental journey" is somewhat less entrenched. Accordingly, "mental journey" feels a little less literal. The phrase "ethnic cleansing" uses the same grammatical pattern, but the conceptual connections it evokes are much less entrenched, and the vocabulary is less

entrenched. It was judged to be highly figurative when first used, but the effect seems to be wearing off with frequent exposure.

Some connections evoked by "figurative" examples might interfere minimally with our category connections and thus be easily assimilated. For example, "A leopard is a tiger with spots instead of stripes" is of course "literally" false and moreover calls explicitly for mental blending that may strike us as figurative. But the connections that we do construct for this expression come to sit nicely inside our category structure for large mammals, partly because they help us to extend that category structure in ways that do not disrupt it. In contrast, connections evoked by some other "figurative" expressions might be deeply disruptive, with the consequence that their assimilation will be resisted by the conceptual apparatus we already have in place. A surprising expression like "time is the whiteness of the wave," which leads us to form weird conceptual connections that challenge our category structures, may not settle readily into our conventional knowledge. It may remain suggestive, never achieving a stable location. It may not be used up—assimilated and naturalized—as we go through it repeatedly: we may be able to return to it again and again, and find it fresh, even powerful, because the connections it suggests cannot be established in our category structures (or perhaps even in our conventional conceptual apparatus) with impunity.

In summary, in my 1989 analysis of the literal versus figurative distinction, I proposed that we feel products to be "literal" or "figurative," that these products arise from the identical cognitive and linguistic mechanisms; but that they evoke different reactions depending on (1) the relative status and degree of entrenchment of the relevant mental arrays in the conceptual structures brought to bear on them, and (2) the degree of entrenchment of the language used for evoking those connections.

Recently, Gilles Fauconnier and I have jointly developed a model of conceptual connection that generalizes and extends my earlier view that "literal" versus "figurative" does not refer to a difference in basic cognitive operations. In the next few pages, I sketch the principles of our model. Subsequently, I draw its implications for the literal versus figurative dichotomy.

In Fauconnier & Turner (1994, 1996, 1998a, 1998b, and 2002), Turner and Fauconnier (1995), Fauconnier (1997), and Turner (1996a and 1996b), Gilles Fauconnier and I have presented our "network model of conceptual integration." The model has additionally served as the basis for Coulson (1995, 1997), Freeman (1997), Mandelblit (1995, 1997), Oakley (1995), Ramey (1997), Sun (1994), Veale (1996), and Zbikowski (1996).[i] The following presentation borrows from these publications.

Conceptual integration is a basic cognitive operation that operates on two input mental spaces to yield a third space, the *blend*. For example, in "Vanity is the quicksand of reason," one input space has quicksand while the other has vanity and reason; the blend has traps for reason.

---

[i]The website for conceptual integration is at http://blending.stanford.edu

In blending, there is a partial cross-space mapping between the input spaces. In the *quicksand* example, the *traveler* in one input is the counterpart of *reason* in the other input.

Perhaps needless to say, the content of the input spaces depends on the domains and frames from which they are built. In the movie *Lawrence of Arabia*, there is a scene in the dry heart of the burning desert where quicksand swallows a child whole. Most people I know have this "dry quicksand frame" available. Others have only a scientific frame for quicksand, in which it is a combination of sand and water and occurs only where the water table is high, making the scene in *Lawrence of Arabia* impossible. I choose the "dry quicksand" interpretation to work with.

Partial structure is projected to the blend from the input spaces. The *quicksand of reason* blend has, from the *quicksand* input, a dangerous trap, namely quicksand, but the blend does not take from the *quicksand* input the knowledge that travelers typically avoid deserts except in the rare instance when they must cross them. From the *reason* input, the blend takes noble and valuable effort but not (for example) vanity as a spur to honorable achievement.

The blend has emergent structure not provided by the inputs. In the *quicksand of reason* blend, the traveler can be ignorant of the trap even when he is in it.

In addition to the inputs and the blend, conceptual integration involves a *generic space.* The generic space contains the structure that is taken at any moment in the construction of the network as applying to the two input spaces. For example, the generic space for "vanity is the quicksand of reason" has action (not specified as physical or mental) intended to achieve something and a difficulty for that action.

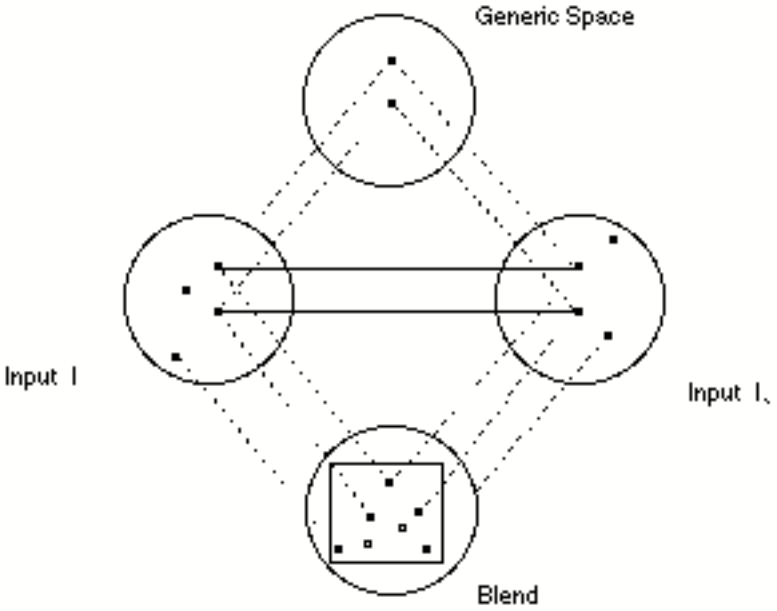

In Fauconnier and Turner (1998b), we present a taxonomy of types of integration networks that arise often. This taxonomy depends upon the notion of

an *organizing frame,* a frame that specifies the nature of the relevant activity, events, and participants. Examples of organizing frames are *man walking along a mountain path*, *boat sailing along an ocean course*, and *gunslingers at high noon.*

The first type of integration network is a *mirror network*, in which all spaces—inputs, generic, and blend—share topology given by an organizing frame. Two of our standard examples of *mirror networks* are "Debate with Kant" and "Regatta." In "Debate with Kant" (Fauconnier and Turner, 1996), a modern philosopher running a seminar says something like, "I claim that reason is an autocatalytic somatic complex adaptive system that develops in the individual. Kant disagrees with me on this point. He says it's innate, but I answer that that's begging the question." In one input, Kant is thinking and musing and perhaps writing. In the other input, the modern philosopher is thinking and communicating; the generic space has a philosopher working on a question; the blend has Kant and the modern philosopher; moreover, the blend has them debating. All of these spaces have the organizing frame, *philosopher considering a philosophical problem.* The blend has an extension of this organizing frame: *two philosophers considering a philosophical problem and moreover debating it.*

In another mirror network, "Regatta" (Fauconnier & Turner 1994, Turner & Fauconnier 1995), a freight-laden clipper ship, *Northern Light*, set the record for an ocean voyage from San Francisco to Boston in 1853 and a modern catamaran is in the process of making that run in 1993. *Latitude 38* reports, "As we went to press, Rich Wilson and Bill Biewenga [the crew of the catamaran] were barely maintaining a 4.5 day lead over the ghost of the clipper *Northern Light*."[ii] Here, all four spaces have the organizing frame *boat making an ocean voyage*. The blend has an extension of that frame: *two boats making ocean voyages and moreover racing as they make them.*

A simpler type of integration network is a *simplex* network, in which one input is a familiar abstract frame and the other input is a relatively specific situation. If we wish to say that two people—John and James—stand in a certain kin relation, we say something like "John is the father of James." The frame of kin relation is one input; the other input has John and James. In the blend, John is the father of James, and there is a role *father of James.*

A conceptual integration network is *single-scope* if the inputs have different organizing frames and only one of those frames is projected to organize the blend. For example, a cartoon of Bob Dole and Bill Clinton having a shoot-out evokes a single-scope network: the frame *gunslingers at a shoot-out* is projected from one of the inputs to organize the blend. The network is therefore *single-scope*.

Any particular simple metaphoric single-scope network—like the shootout between Dole and Clinton—may have inhering within it a higher-order conventional metaphoric mapping, called by Lakoff and Johnson (1980) a *basic metaphor*. In the case of the *shootout* network, the inhering basic metaphor is *opposition is combat.* Such a basic metaphor is highly productive and inheres in

[ii]"Great America II," *Latitude 38*, volume 190, April 1993, page 100.

many particular constructions of meaning but is itself abstract. It never constitutes an active, complete, on-line construction of meaning. It always requires additional conceptual specification and projection.

A conceptual integration network is *double-scope* if the inputs are organized by different frames and some topology is projected from both input frames to organize the governing frame of the blend. The metaphor "vanity is the quicksand of reason" is a double-scope network with frame structure projected from both inputs to organize the blend. The projections from the organizing frame of the *quicksand* input are obvious: the blend has a traveler, a path traveled, distance traveled, motion, a potential trap that arrests motion, and so on.

But frame-level projections come from the *reason* input, as well. Consider first intentional structure: the reasoner can be unaware of his failure even when his failure is nearly complete. This is projected to the blend, in which the traveler/reason can be unaware of being in quicksand. The traveler/reason can be deluded, viewing himself as perfectly rational, oblivious to the fact that he has in fact long been trapped. This intentional structure conflicts with the frame of the *quicksand* input, in which it is unconventional to be ignorant that one is in quicksand, unconventional to think that one is traveling normally when one's torso is sinking.

Next, consider causal structure from the *reason* input: reasoning can lead to vanity about one's reasoning, which can lead in turn to diminished reason. This structure projects to help organize the blend: in the blend, quicksand/vanity exists for the reasoner but not for the person whose mind is merely wandering, even though they are both travelers. This causal structure conflicts with the organizing frame of the *quicksand* example, in which traveling is not causally related to the existence of quicksand, and in which all travelers in the desert face the same dangers. Additionally, in the *reason* input, the more you have achieved through reason, the more justification you have for being vain; in the blend, the more you have achieved through reason, the more vulnerable you are to being caught in quicksand. But this structure conflicts with the *quicksand* input, where it is novice travelers who should be most vulnerable to quicksand.

Now consider the structure of roles in the *reason* input: there is only one reasoning capacity. The blend follows this structure: the traveler is solitary, or, if not solitary, then accompanied by *unequal* companions (character, memory, etc.). This structure of roles conflicts with the *quicksand* input, where there may be several equal travelers.

Now consider modal structure from the *reason* input: the reasoner does not have the choice of foregoing reasoning while remaining intellectually sophisticated. This projects to the blend: the traveler cannot choose to forego traveling in deserts; traveling/reasoning always presents a certain danger; that danger is in the desert exclusively; so the traveler/reasoner must deal with the desert. This structure conflicts with the *quicksand* input, in which the traveler can avoid the danger by declining to travel through deserts (which can be viewed as uninteresting in any event)—there are many wonderful places one can visit as a

sophisticated traveler; one can experience a lifetime of interesting travel without entering a desert; and so on.

In summary, although the frame-level projections to the blend from the *quicksand* input are obvious, there are frame-level projections of intentional, causal, modal, and role structure from the *reason* input to help organize the blend, and these projections conflict with the frame of the *quicksand* input. The blend is in these ways double-scope.

What are the implications of the network model for the literal versus figurative distinction? The network model generalizes my earlier claim that the same conceptual and linguistic operations underlie "figurative" and "literal" examples. Different examples will feel "literal" or "figurative" for a number of reasons, including type of network. The type of the network has to do with the relative status of counterparts in the cross-space mapping between the inputs, a status that is judged with respect to the category structures and related conceptual structures brought to bear on them.

For example, a simplex network like "John is the father of James" has two inputs with the following relative status: one is a familiar abstract frame, while the other is a relatively specific situation with no competing frame. The familiar abstract frame is routinely applied to the conceptual domain (individual human beings) upon which the specific situation is built. This type of integration network usually feels highly literal.

By contrast, if two inputs come from what are felt to be widely different specific conceptual domains, the result is a different type of integration network, either single-scope or double-scope. The structure that applies to both of them (i.e., the generic space) is typically highly abstract relative to both of the inputs. Such a case is commonly felt to be figurative (depending, as we will see, on some other gradients of distinction). In particular, highly double-scope networks, like *Vanity is the quicksand of reason*, are typically judged to be highly figurative.

It turns out that even this taxonomy of types of integration networks is too rigid: distinctions between the types are in fact graded, and judgments of literal versus figurative are accordingly graded. Let us consider some examples of grading, taken from *Death is the Mother of Beauty* (1987) and further analyzed in Fauconnier & Turner (2002).

As we have seen, "John is the father of James" seems fully literal; there is no competition between organizing frames of the inputs, and the kinship frame is routinely applied to the conceptual domain of individual human beings. "Zeus is the father of Sarpedon"—where Sarpedon is the mortal son of Zeus by a human woman—may strike us as less literal because the kinship frame meets some resistance from the Sarpedon space and the integration is slightly double-scope: from the Sarpedon input, the blend receives the immortality of Zeus; from the kinship input, the blend receives the ego-father relationship but cannot receive the mortality of the father. Yet the role *mother* in the kinship frame does have a standard counterpart in the Sarpedon space, as do various stages of human progeneration involving the mother, and these counterparts are fused in the blend.

A slightly different case is "Zeus is the father of Athena." In Fauconnier & Turner (2002), we write, "The blend does not take the frame-level structure *sexual act with a woman that leads to conception and progeneration of an infant*. It takes something more general: the causal link between the parent and the existence of the offspring (although not the immaturity of the offspring), the emergence of the offspring from a container-like body part of the parent, paternal responsibility and protection, and inheritance of attributes."

Let us consider an extended example of double-scope blending—Milton's portrayal of Satan as father in the second book of *Paradise Lost.* I analyze this passage in Turner (1987). The commonplace notion of Satan is already a blend for which a conceptual domain has been elaborated. Satan is a blend of individual human being—thinking, talking, desiring, intending, and so on—and theological ontology. In the theological space, there are eternal features (e.g. evil) as well as non-human powers and limitations. Satan is anthropomorphic, but he has theological features and unhuman conditions. The blended domain for Satan is quite elaborated—Satan has like-minded colleagues in the form of a cohort of devils; Satan and the devils form an intricate hierarchical organization of social groups; and so on. This blended domain is entrenched both conceptually and linguistically. Consequently, although the blend is in some ways double-scope, expressions like "The devil made me do it" or "Get thee behind me, Satan"—or even expressions based on further blending, such as the reference to a child as a "little devil"—do not feel especially figurative.

Milton recruits new structure to the inputs. His purpose is to develop a blend with yet further emergent structure. The result is an integration network that is less entrenched both conceptually and linguistically and that is aggressively, explicitly, and idiosyncratically double-scope. It accordingly strikes us as thoroughly figurative.

Milton activates for the theological space evil, disobedience, sin, death, and their relations, as well as the psychology of the prototypical sinner confronted with spiritual death. He activates for the human space progeneration and kinship relations, especially the role *father.* He adds to the human kinship space a pre-existing blend, of the birth of Athena from the brow of Zeus. In Milton's blend, Satan conceives of the concept of sin; a fully grown woman, Sin, leaps from his brow. Satan is attracted to sin/Sin: he has sex with her. Although he does not know it at the time, his involvement with sin/Sin has a consequence, namely death—in the blend, Death is the male offspring of Satan's incestuous involvement with Sin. Death rapes his mother, causing her to give birth to a small litter of allegorical monsters.

After Satan has been sent to Hell and has decided to try to escape, he meets two characters at the gates of Hell who have been stationed there to keep him in. They are Sin and Death. He does not recognize them.

I explain in Turner (1987) how the two input spaces—the human space and the theological space—correspond in some ways but not others. Milton chooses to draw from one or the other as it suits his conceptual purposes. In the new vocabulary of the network model, my earlier discussion of Milton's passage analyzes it as a selective, double-scope projection to a blended space. For

example, Milton takes from the space of human beings and kin relations Sin's intercession between Death and Satan—father and son—when they are on the brink of terrible combat. By contrast, he takes exclusively from the theological space many central features, as follows. In the theological space, there is a sinful cast of mind which does wish to understand that mortality and spiritual death are the consequences of sin and which at last becomes appalled when it is forced to acknowledge these consequences. Hence, in the blend, Sin is surprised to have conceived Death, and she finds her son odious. Next, in the theological space, mortality and spiritual death overshadow the appeal of sin and are stronger than sin; acknowledging death devalues sin; willful, sinful desires are powerless to stop this devaluation. Hence, in the blend, Sin is powerless to stop her horrible rape by Death. In the theological space, the fact of spiritual death brings ceaseless remorse and anguish to the sinful mind, and the torments of hell bring eternal punishment. Hence, in the blend, the rape of Sin by Death produces monstrous offspring whose birth, life, actions, and relationship to their mother are impossible for the domain of human kinship:

> These yelling Monsters that with ceasless cry
> Surround me, as thou saw'st, hourly conceiv'd
> And hourly born, with sorrow infinite
> To me, for when they list, into the womb
> That bred them they return, and howl and gnaw
> My Bowels, thir repast; then bursting forth
> Afresh with conscious terrors vex me round,
> That rest or intermission none I find.

We see here Milton's skill as a blender. When he takes structure from one input, he is adept at seeking out suitable structure to recruit to the other input, so that the two structures can be given counterpart relations and blended. Children are not prototypically disliked, but Milton can recruit the unusual scenario of disliking a child so he can blend it with horror at recognizing the fact of death. Sons do not typically rape their mothers, but Milton can recruit that horrible scenario so he can blend it with death's effect on sin.

Milton's ingenuity as a blender is best shown, I think, in his recruitment of a particularly vivid medical frame to the input of human kinship. This medical frame is traumatic vaginal birth that physically deforms the mother. In the human space, this disfiguration makes the mother subsequently less attractive. Milton places this newly recruited structure into counterpart relation with something crucial in the theological input—the fact that sin becomes less attractive when death appears as its outcome. The blend is particularly grim:

> At last this odious offspring whom thou seest
> Thine own begotten, breaking violent way
> Tore through my entrails, that with fear and pain
> Distorted, all my nether shape thus grew
> Transform'd.

My original analysis of Milton's portrayal of Satan as father provided an inventory of its elaborate selective projection, emergent structure, double-scope structure, multiple blending, dynamic recruitment to the inputs of additional structure, maintenance of connections to the inputs, and projection of inferences back to the inputs. But it did not use these terms and did not connect Milton's passage systematically to the many examples of blending in other domains of human thought, language, and action. Less narrowly I presented in *Reading Minds* an analysis of XYZ constructions (like "Vanity is the quicksand of reason") as involving a basic mapping scheme that invokes open-ended conceptual work that leads to emergent structure. I inventoried examples of cross-space mapping, selective projection, and emergent structure, but my analysis of these cognitive operations was incomplete, and my assertion of a broad scope for XYZ mappings was—surprisingly—too modest by far.

These earlier analyses are subsumed by the newer Fauconnier and Turner network model, which gives a much fuller analysis of the cognitive operations involved in conceptual projection, a specification of taxonomies of types of integration networks, a set of optimality constraints on creating them, and a program for demonstrating the general scope of conceptual integration. We have now connected my kinship metaphor and XYZ examples to examples that look ostensibly altogether different—the invention of complex numbers, the operation of grammatical constructions, the evolution of syntax, action slips, category extension, counterfactual argument, and so on.

Although Milton's portrayal of Satan as a father is double-scope, it preserves considerable structure associated with *father* and *birth*. Consider first the paternity of Death. The "father" has human form and speaks human language, is excited by feminine beauty, and has anthropomorphic sex with an anthropomorphic female in a prototypical human scene. There is a birth through a vaginal canal. The son inherits attributes of both father and mother. Father and adolescent son have a conflict over authority. And so on. Now consider the paternity of Sin. The father again has human form and speaks human language. There is an offspring in human form, who emerges from a container-like body part and who develops into a sexual being.

Other examples, taken from *Death is the Mother of Beauty*, show a different projection from the space of *father* and *birth*. "Satan, liar and father of lies" does not take the anthropomorphic offspring. "The acorn is the father of the oak" takes neither anthropomorphic form nor anthropomorphic progeneration for either father or child. "Thy wish was father to that thought" (Shakespeare) does not take physical distinction for either father or child. Similar double-scope structure appears in "Fear, father of cruelty" (Ezra Pound), "Pain is the father of complaint" (Sidney), "Love's extremity is the father of foul jealousy" (Spenser), and "Pale desire, father of Curiosity," (Blake).

Consider as a final example the XYZ expression "The Child is Father of the Man" (Wordsworth). The two inputs—father-and-child versus child-growing-to-man—come from the same conceptual domain, human life. But the example feels figurative, for the following reasons. First, the cross-space connections are

highly resisted because they run counter to usual categories: *Immature child* in the first input has as its counterpart *father* in the second input, and *grown man* in the first input has as its counterpart *immature child* in the second input. Second, the blend must integrate frame-level structure from both inputs in a particularly surprising way. The chronological *child* in the blend takes from the input of father-and-child the relative influence (and even causal role) of the father, but it takes from the input of child-to-grown-man the relative *youth* of the child. The chronological *man* in the blend takes from the input of child-to-man the maturity of the man, but it takes from the input of father-and-child the dependency of the child.

The oddness of its counterpart connections and the extensive double-scope structure of its blend help make Wordsworth's line feel figurative. But the syntax and mapping scheme of "The Child is Father of the Man" are the same as the syntax and the mapping scheme of "John is the father of Mary." Both evoke a conceptual mapping scheme involving conceptual blending, but "John is the father of Mary" seems absolutely literal.

"Peeled apple" also seems absolutely literal. But as Gagné and Murphy (1996) write:

> Understanding a combined concept involves creating a new concept. For example, a *peeled apple* is no longer just an apple—its features are not entirely identical to those of an apple. A peeled apple is white, not red, and a peeled apple is more likely to be used for baking than is an unmodified apple, and so on. In short, the concept of the head noun is modified in some way by the addition of the modifier. Although one might think that this modification would be a simple process of adding the meaning of the modifier to that of the head noun, this has not turned out to be the case. The interpretation of combined concepts involves an interaction between the two constituents, rather than an additive process. For example, the fact that peeled apples are white is not part of the meaning of *peeled*, but is inferred, based on our interpretation of the entire phrase. Peeled oranges are not white, are not likely to be used in cooking, and so on. Thus, *peeled* cannot be adding the same feature to *apple* and *orange*. (Gagné and Murphy, 1996, 80).

From the perspective of the network model, "peeled apple" evokes a simplex network. One input has a general frame of *peeling* and the other input has *apple*. The two words are prompts for putting together two provisional input spaces. The blend has considerable emergent structure—such as whiteness and association with baking pies—that is not given for the inputs. There is a cross-space mapping connecting, for example, the apple to the object being peeled. There is selective projection—we do not project the color of the apple, or peeling with the fingernails, or peeling as a natural process (*do not apply ointment to peeled skin)*, and so on. Completion occurs in the blend through recruiting the frame of

baking (for example). We show in Turner & Fauconnier (1995) that an example like "peeled apple" is not unusual. Conceptual integration can be seen not only in striking examples like "land yacht," "jail bait," and "Chunnel," but also in unremarkable examples like "waterproof," "tamper-proof," "foolproof," "child-proof," "talent pool," "gene pool," "water pool," "football pool," "betting pool," "door knob," "radio knob," "house boat," "boat house," and "black bird."

"Fire station," for example, feels entirely literal. But a fire station does not have fire, provide fire, or receive fire; fire is not part of *station* or the category that includes *station*. We have a mental space with *fire* and a mental space with people or equipment stationed at a *station* for a purpose, and we can integrate these mental spaces conceptually into a story in which fire is not a feature of the station or a counterpart of the station. In this story, the equipment and people at the station go to deal with fire. *Fire station*, like *peeled apple*, is a simplex network: the frame of stationing equipment and agents to deal with something is applied to the input *fire*. "Fire station" asks us to create this simplex network. It does so by means of a highly entrenched phrase learned early in childhood. The result is a conventional integration that sits easily in category structures, since we are familiar with categorizing by purpose.

Milton's infernal trinity, *peeled apple,* and *fire station* arise from the same cognitive operation—conceptual integration—but the infernal trinity seems highly figurative while *peeled apple* and *fire station* seem absolutely literal. The counterpart connections in *peeled apple* can be accommodated in our category structures: we already have a way of seeing transformations of objects as categorical subtypes of the object (shrivelled apple, rotten apple, etc.). The application of the frame of peeling to the domain of fruits and vegetables is highly frequent and familiar. The phrase "peeled apple" is entrenched both as a pattern ("stewed carrots," "minced onions," etc.) and as a specific item.

Similarly, the counterpart connections in *fire station* can be accommodated in our category structures—the station is set up to deal with *something*, and that *something* has as its counterpart *fire*. In the blend, there is indeed a fire, and the agents and equipment at the station perform the action of dealing with it. This blend may be entirely counterfactual—imagine a fire station as a precautionary element in a chemical plant where no fire ever erupts during the entire existence of the plant—yet the station will be no less fully a "fire station." The simplex integration network feels entirely literal because we already categorize stations according to what they are designed to deal with, because we routinely apply the frame of *station* to the domain of *fire* and to the super-domain of *crisis* or *disaster*, and because the conceptual connections and linguistic forms in "fire station" are entrenched.

By contrast, Milton's infernal trinity is a different type of integration network, highly and aggressively double-scope, explicitly novel in much of its conceptual structure and its linguistic expression. Our judgments of the packages differ, but the basic cognitive operations used to construct them do not.

Let us consider a final suite of connected examples that may help tease apart several aspects of the literal versus figurative distinction. The sentence "President Franklin Delano Roosevelt moved at a quick pace during his first one

hundred days in office" seems essentially literal to many people. Yet it is grounded in a conceptual blend. One input has Roosevelt's achievements; the other has a person moving along a spatial path toward destinations. In the cross-space mapping, the traveler is the counterpart of FDR. The generic space taken as applying to both inputs has an abstract agent, abstract purposive actions, and an abstract oriented linear scale whose locations correspond to grades of achievement of those purposes. In the blend, the linear scale that is the spatial path of one input is fused with the linear scale for measuring achievement in the other input. In the blend, to be farther along the path is to have accomplished more of the relevant purposes. This is a single-scope network: the frame of the blend is an elaboration of the organizing frame of only one of the inputs, the *travel* input.

Although this network of FDR-as-runner connects two quite different conceptual domains, it can feel literal, for various reasons. First, the basic network of which this is an instance is highly entrenched. It forms what Lakoff and Johnson (1980) call a *basic metaphor*. It is constructed repeatedly in indefinitely many cases that differ only in the specific details of the target input and blend. Exactly this input of motion along a path toward destinations and this generic space are projected in exactly this way to many target inputs whose organizing frame is purposive activity. The result in all these cases is exactly this blend, not counting specific details. New on-line construction of meaning in this case is limited to specific details like the identity of the agent (FDR), the particular kind of purposive activity (legislation, government), and the interval of time (one hundred days).

The generic space for this network (agent with purposes and a linear scale of success) is moreover entrenched in its own right, accessible for projection to any purposive activity. Indeed, that generic structure is now entrenched in the frames of various purposive activities themselves, which carry the vestiges of the conceptual integration networks in which they are embedded. In such cases, we do not need to activate the entire network fully and we do not need to perform on-line invention of new projections at the frame level.

This integration network (*purposive agent as traveler on a path*) is moreover our standard cognitive instrument for thinking about purposive activity, and it is used with very high frequency.

Finally, the vocabulary "move at a quick pace" has historically been projected to the generic space, the conventional frame of the target, and the conventional frame of the blend, and is entrenched there.

We can alter the example, first by using vocabulary more tightly tied to the source input: "FDR *made the dust fly as he sped along* during his first one hundred days." Next, we can choose vocabulary that evokes a particular scenario for the source: "FDR moved at *full gallop* through his first one hundred days." In these cases, we must make the minor but indispensable inference that someone who speeds along while making dust fly or who moves at full gallop in fact moves a far distance over the path.

Next, we can point explicitly to an additional scenario and highlight the existence of a counterfactual blend, as in "If FDR had been a sprinter, he would

have won the Olympic gold for his performance during his first 100 days in office." This now feels thoroughly figurative. The type of integration network is unchanged, as are the basic cognitive operations involved, but we have changed the degree of entrenchment of the vocabulary, the amount of on-line blending needed, the familiarity of the scenario as applied to this purposive activity, and the degree of explicit acknowledgment of the blend.

We can also imagine a second and nearly identical conventional blend, "President Clinton has moved at a slow pace during his first one hundred days in office." We can make a comparison between these two conventional blends: "F.D.R. moved at a quick pace during his first one hundred days; President Clinton by comparison has not." This sets up a counterpart mapping between the two specific blends of FDR-as-runner and Clinton-as-runner. The counterpart mapping connects president/runner to president/runner, FDR to Clinton, and so on. These two blends are both specifications of the more abstract conventional blend *purposive agent as traveler on a path.*

These two blends—FDR-as-runner and Clinton-as-runner—can themselves be input spaces to a new, hyper-blended space, as when we say, two months after President Clinton has taken office, "Clinton was supposed to hit the ground running. He implied that he was going to accomplish as much in his first 100 days as F.D.R. accomplished in his. So far, Clinton has failed completely *to keep pace with* F.D.R." This is a *mirror network:* the organizing frame shared by the two (already blended) inputs, their generic space, and their hyper-blended space is the already blended frame *American president as traveler on a path.* In the hyper-blend, which has both FDR and Clinton, this already blended frame is extended. First, it is extended through composition: although the two paths have been projected from the two inputs to a single fused path in the hyper-blend, the two agents from the two inputs are projected to discrete agents on that single path, so that now we have not one president/runner on the path but two. Second, the blend is additionally extended through completion: the frame of a *race* is used to complete the blend. It brings with it the structure of *keeping pace with*, *being ahead or behind*, and so on, which is emergent structure unavailable from the inputs themselves.

Although the shared frame of *American president as traveler on a path* is fairly conventional, emergent details of the blend are emphasized ("keep pace with"). This hyper-blend can be made to feel increasingly figurative the more on-line work we require, the less entrenched we make the vocabulary, and the more attention we draw to the blend, as in "At this rate, Clinton's term will be over before he gets anywhere near *the finish line*." Here, we point directly to the frame of *race*, which is in the hyper-blend. To construct this finish-line hyper-blend, we must do considerable on-line work to conceive of a finish line that corresponds to F.D.R.'s degree of accomplishment on his hundredth day in office in the relevant input space of *F.D.R.'s first year in office*. Finally, we can guarantee that the blend is forced into consciousness and is felt to be figurative by requiring the construction of a provisional conceptual domain, as in "Clinton is in a race with the *ghost* of F.D.R."

In all of these cases, the conceptual networks are formed using the same cognitive operations. The results feel more or less literal or figurative for various reason, but not because they have been formed through fundamentally different cognitive operations.

Two traditional questions taken as guiding inquiry into the literal versus figurative distinction are, "Is figurative thought mirrored in figurative language?" and "Is figurative thought necessarily paired with figurative linguistic form?" Under the present account, these questions are misdirected.

Typically, languages already possess constructions that can be used to evoke any sort of integration network. "Boat house," "jail house," and "door knob" use compound nouns and existing lexical items to evoke conceptual integrations. So do "land yacht," "fossil poetry," and "jail bait," which seem "figurative." "He kicked the ball over the fence" uses existing lexical items in the existing Caused Motion Construction (Goldberg 1995, Fauconnier & Turner 1996), to evoke the blending of (1) a set of unintegrated actions and events (he kicked, his kick made contact with the ball, the ball moved, the trajectory of the ball's motion was over the fence) with (2) the already integrated but abstract Caused-Motion story, in which an agent's action causes an object to move in a direction. I found the following "figurative" example in the *New York Times:* "So far, the people of this small textile town in northwestern Carolina have been unable to pray Mrs. Smith's two little boys home again." This "figurative" example equally uses existing lexical items in the Caused-Motion construction to evoke a similar blend. The cognitive and linguistic operations are the same in the two cases. What varies between them is instead the relative category status of the inputs and the familiarity of applying the Caused-Motion frame to the other domain (*body actions* versus *praying*). We rarely if ever use the intransitive verb "pray" in the three-argument Caused-Motion construction, but when we do, the linguistic operations are no different from those used in "He kicked the ball over the fence." In Fauconnier & Turner (1996), we analyze a sequence of similar Caused-Motion examples that fall at various points on the gradient of the literal versus figurative distinction: In "Junior sped the car around the Christmas tree," "sped" evokes the motion of the object; In "Paul trotted the stroller around the park," "trot" evokes the action of the agent; In "Sarge let the tanks into the compound," "let" evokes causal connection; In "Max carted the drums into the warehouse," "cart" evokes a vehicle used; In "Jane muscled the boxes over the fence," "muscle" evokes the part of the body used for the action; In "The spy Houdinied the drums out of the compound," "Houdinied" evokes someone associated with actions of a certain character. There are no new linguistic constructions in any of these examples, although some feel figurative.

Rarely, new grammar or lexical items do arise under pressure from conceptual blending. In Turner and Fauconnier (1995), we report that page one of *The Atlanta Constitution* of 17 February 1994 carried a header reading "Out on a Limbaugh," followed by a summary of the story on the inside pages: "Critics put the squeeze on Florida's citrus industry for its $1 million deal with broadcaster Rush Limbaugh." The formal blend of "out on a limb" with "Limbaugh" is driven by a conceptual blend of (1) an agent who climbs out on a limb of a tree with (2)

the deal between the Florida citrus industry and conservative radio talk show host Rush Limbaugh. It turns out that conceptual counterparts that are conceptually blended (*limb, Limbaugh*) have formal expressions that are formally blended ("limb," "Limbaugh"). There is emergent formal structure in the blend as follows. "Out on a limb" has an indefinite article with a common noun. "Limbaugh" is a proper surname, not a common noun. Although a proper surname in English can become a common noun for a group of people with that surname ("She's a Kennedy," "She's the poorest Kennedy") or a group of people analogically equivalent to a particular person of that surname ("He's an Einstein"), here "Limbaugh" is not used as a common noun, referring to namesakes or analogs of Limbaugh. Yet it follows an indefinite article. Following an indefinite article is a property of its counterpart formal element, "limb," associated with the other input to the blend. The blend has a new formal element consisting of previously unavailable syntactic structure—indefinite article + proper name.

We often feel that new and deviant language is "figurative"—indeed, "Out on a Limbaugh" is a prototype of a figurative pun. It seems to ask for laughter. But now consider the following example. At the 1988 Olympics in Korea, a boxing match between an Australian contestant and a Korean contestant ended with strange events, including a skirmish that involved officials and coaches. The Australian coach, interviewed at the airport before boarding a plane to leave in disgust, said, as closely as I can recall, "I was hit by the judge; I was tried to be hit by the umpire." We count the second verb phrase as a mistake, but it is not an arbitrary mistake. It follows principles of formal blending under pressure from conceptual blending, as follows. The speaker has one input space in which he is the victim or patient of actions. That scene comes with useful syntax, namely the Passive Construction ("I was bit," "I was made to cry," "I was insulted"). The speaker has this scene and this syntax active and wishes to perpetuate them in the minds of the members of his audience. He also has active the set of unintegrated events in which the umpire is an actor and he, the Australian coach, is the umpire's victim or intended victim. In this scene, the umpire tries to do something, and what he tries to do is hit the coach. Had the coach located the verb "assault" as language for this scene, he could have continued to use the Passive Construction with perfectly grammatical parallelism: "I was hit by the judge; I was assaulted by the umpire." But either he did not locate "assaulted" or "assaulted" seemed wrong for some reason such as inappropriate register or lack of viscerality and vividness. There is other syntax available for this scene, in which the verb phrase is active—"The umpire tried to hit me"—but it does not evoke so clearly the established abstract scene of passivity that the Australian coach wishes to keep active. The coach wants to prompt for that scene by using the Passive Construction, but he cannot use the syntax of "try" and "hit" in the Passive Construction because "try" as an auxiliary verb does not take the passive form. The coach therefore creates a formal blend—*try* as an auxiliary that takes the passive form—in order to express the conceptual blend. He may have received additional help in constructing this formal blend from existing syntax in expressions like "This tool was designed to be used by the designer," where the

designer is the agent of both the action of the designing and the action of using, just as the umpire is the agent of both the action of trying and the (unachieved) action of hitting. "Design" passivizes while "try" as an auxiliary verb does not, but the Australian coach leaves behind that part of the syntactic structure as he gives "try" a new, emergent syntax under pressure from conceptual integration: "I was tried to be hit by the umpire."

"Out on a Limbaugh" and "I was tried to be hit by the umpire" use the same operations of conceptual and linguistic blending, but the first feels "figurative" and the second feels like a mistake. In "Out on a Limbaugh," at the conceptual level, we have a single-scope network whose generic space is abstract relative to the organizing frames of the inputs. In "I was tried to be hit," we have something close to a simplex network, in which a frame of passive victimization is applied to physical actions; moreover, the application of this frame to this conceptual domain is archetypal and routine. The emergent syntax of "Out on a Limbaugh" seems to be planned, while the emergent syntax of "I was tried to be hit" seems to have arisen spontaneously. For reasons such as these, the first feels figurative and the second does not, and the first seems witty and the second seems a little embarrassing. But the basic cognitive and linguistic operations are not different.

Another traditional question taken as guiding inquiry into figurative language and thought is, "How do figurative thought and language evolve?" The short answer is, conceptions and forms that feel figurative evolve in the ways that all thought and language evolve; some products in that evolution feel more or less "figurative" according to their location on the interacting gradients of distinction, but this feeling will vary across persons and moreover does not indicate a fundamental difference of cognitive operation.

A substantive answer to this question would be very long because it would be a theory of the evolution of conceptual structures and linguistic forms. Such a theory would be highly complicated since human thought and language arise through the interaction of several complex adaptive systems, including biota (all living things through all time; a unit is a gene pool and all its ancestor gene pools); a given gene pool (a unit is a gene); all conceptual systems in all individuals over all time; a conceptual system shared by a community, and all the conceptual systems that are ancestors of that conceptual system; a conceptual system within a single individual, and all the conceptual systems that were, in the individual, ancestors of the current conceptual system; human language, all of it, over all historical time; a human language shared by a linguistic community and all the diachronic linguistic structures that are ancestors of that language; and a human language, in an individual, and all the linguistic systems that were, in the individual, ancestors of that current linguistic system.

This list, already paralyzing in its complexity, is actually more complex, since its elements overlap and interact. Modeling thought and language (and therefore thought and language that feel "figurative") involves modeling its interacting complex adaptive systems. The network model is only a modest gesture in this direction. In it, existing conceptual and formal elements and their pairings are inputs to an operation of integration that is selective and that results

in emergent structure. Outputs of integration can become inputs to integration. The result is pathwise development of a system in which elements stand in relation to other elements. What can arise in the system at any moment in its evolution depends on what has already arisen that survives. The system is dynamic; it never stands still. Conceptual integration exploits accidents as a fundamental part of its functioning; indeed, basic ("literal") structure in the system can arise from the exploitation of remarkable accidents. Products of integration that feel at one time figurative may feel at other times literal. Formal blending to create new forms may be guided by pressure from conceptual blending. These operations are not deterministic or algorithmic, but instead are guided by optimality principles and by degree of success in the moment of operation. In my view, the cognitive operations involved in the evolution of the conceptual and formal patterns we see in "figurative" examples like "land yacht" or "jail bait" are identical to those we see in "literal" examples like "fire station" or "brown cow."

A last traditional question taken as guiding inquiry into figurative language and thought is, "What is the appropriate relation of an abstract theory of figures to a rich theory of individual figurative events?" Actual figures occur only in dynamic, on-line construction of complete meanings. The study of figure typically does not focus on this condition. The central products of the study of figure are typically lists of abstract elements—"figures"—with examples: here is antithesis and here are examples; here is metonymy and here are examples; here is the basic metaphor LIFE IS A JOURNEY and here are examples.

In this style of analysis, the examples are adduced to *refine the elements of the theory* rather than as objects of case study. To define metonymy and provide an example does not supply an analysis of the specific example, or at least the construction of meaning prompted by that specific example. Historically, the study of figure has taken on the job of proposing abstract elements—"figures"—and giving examples, but has not taken on the job of explaining the dynamism and completeness of individual examples. Typically, the study of figure attempts to isolate and exemplify partial structures that get used in the construction of meaning, but not to give a theory of that actual use.

Traditional grammar follows the same pattern: here is a partial structure we call "noun" and here are some examples of nouns; here is a partial structure we call "verb" and here are some examples of verbs; here is a partial structure we call the "passive construction" and here are some examples of passive constructions.

Modern grammar specializes in this kind of analysis of partial instruments: here is verb argument structure, with examples; here is ergativity, with examples; here is inflectional morphology, with examples. None of these abstract partial structures could itself be a full meaning; they are all partial instruments whose utility derives from their availability to be recruited in actual linguistic and conceptual events. The dynamism of the actual full meanings is not modeled. Most models of grammar assume that there exists an abstract object of study—called "the language"—that transcends the full and dynamic

particular linguistic events in individual brains, just as principles of physics transcend actual physical events.

The impulse to construct a theory that consists of abstract elements is strong and understandable, given the success of the mathematical model of theoretical knowledge. Models of mind and language that follow the mathematical tradition look for elemental structures that serve as partial instruments. Theories of semantic primitives, innate concepts, language bioprograms, and symbolic artificial intelligence (such as conceptual dependency diagrams) follow this tradition.

This tradition is not exclusively formal, and it is not exclusively objectivist, either. The theory of basic metaphor (with which I have been associated) attempts to isolate a quite small number of elemental basic metaphors (maybe five or six hundred) that we all know, and to provide examples of each, with the examples meant as evidence for the existence of the abstract elements of the theory.

The central danger for such partial models of conceptual construction is that they might not "scale up" appropriately. The well-known failure of attempts to "scale up" from partial artificial intelligence models to full models is worth remembering in this respect. An analogy from the neurosciences makes the danger clearer: we have a folk theory that we assign color to a spot in the visual field according to the kind of light reflected from that spot in the visual field, but what happens is much more complicated. (Hubel 1995 and Zeki 1993). There are three kinds of cones in the retina, each sensitive to one of three wavebands of light called (inaccurately) red, blue, and green. Suppose we have three projectors, each of which shines one of the wavebands of light, with an intensity we can set on a dial. Suppose we turn on the projectors, at certain settings, to illuminate a painting that consists of rectangles of color. Suppose, finally, we pick out a red rectangle, and measure for each of those wavebands the intensity of light reflected from the red rectangle. Now we look at a green rectangle, and adjust the intensity of light coming from each projector until our measuring device shows that, for each of the wavebands, the identical intensity of light is now being reflected from the second rectangle as was coming from the first when it looked red. *We will still find that the second rectangle looks green and the first rectangle looks red.* The brain is able to compute, for each waveband, a record of differential reflectance of light across the visual field, and then to perform a differential computation across the three differential records, to produce an assignment of constant color under remarkably different conditions of illumination. In this way, we are able to "discount the illuminant" as we attempt to find constancy in the environment. The point of this analogy is that a partial model of color vision does not scale up to a successful model of actual color vision, because what is happening in assigning features to any part of the visual field depends upon the overall activity of vision. We need a model of the operation of the whole in order to account for any part of color vision. Partial models of partial instruments of color vision do not scale up to the kind of global model of computation over global records that is needed to account for color constancy.

The observational data we wish to account for in the case of thought and language all consist of on-line, dynamic construction of full meanings and full expressions. Catalogs of partial instruments that may underlie that data are useful to the extent that they actually help us to account for the data, but it cannot be assumed in principle that the data will be accounted for as linear compositions of individual partial resources. In 1956, George Miller complained that scientific journals had become catalogs of parts for machines the scientists never build.[iii] Cognitive scientists, linguists, and rhetoricians are vulnerable to the analogous observation. The network model is a modest attempt to take a step in the direction of modeling the on-line, dynamic construction of full meanings that arise through conceptual integration. The operation of conceptual integration can recruit from many domains and it can develop elaborate mappings and projections. It is not algorithmic or deterministic, but it is guided by optimality principles that are sensitive to purpose and situation.

From the view of the network model, the contrast of literal versus figurative appears to be unproductive as a theoretical principle for distinguishing cognitive operations. The original view of figure which Quintilian (1921) set aside, in which a figure is any pairing of a formal pattern with a conceptual pattern, and in which figures stand in relational networks, is by contrast basic and indispensable.

Language offers sets of prompts for cognitive operations such as conceptual integration. We conduct those cognitive operations on conceptual structures available to us. A theory of figure that embraces this characterization faces great challenges as it attempts to develop a model of cognitive operations, a model of the relational network of form-meaning pairs that prompt us to perform these cognitive operations, a model of gradients of distinction in the products of those cognitive operations, a model of the ways in which form-meaning pairs arise and evolve, and a model of the ways in which these cognitive operations and figures perform in actual, on-line, dynamic creations of meaning and expression.

The study of figure has been sidetracked from these issues since the Greeks, with the surprising and humbling result that the study of figure, one of the oldest bodies of knowledge in the human sciences, remains in our age still in its infancy.

[iii] As quoted in George A. Cowan, "Conference Opening Remarks," in George A. Cowan, David Pines, and David Meltzer (1994), page 2.

References

Aristotle. *Poetics*. 1995. Edited and translated by Stephen Halliwell. In *Aristotle*, volume 23. Cambridge: Harvard [Loeb].

Aristotle. 1991. *On Rhetoric: A Theory of Civic Discourse.* Translated by George A. Kennedy. New York: Oxford University Press.

Burke, Kenneth. 1969 [1945]. *A Grammar of Motives*. Berkeley: University of California Press.

Burke, Kenneth. 1969 [1950]. *A Rhetoric of Motives*. Berkeley: University of California Press.

[Cicero]. 1954. *Rhetorica ad Herennium*. Translated by H. Caplan. Cambridge: Harvard [Loeb].

Coulson, Seana. 1995. Analogic and metaphoric mapping in blended spaces. *Center for Research in Language Newsletter*, 9:1, 2-12.

Coulson, Seana. 1997. "Semantic Leaps: The role of frame-shifting and conceptual blending in meaning construction." Ph.D. dissertation, UC San Diego.

Cowan, George A., David Pines, and David Meltzer, editors. 1994. *Complexity: Metaphors, Models, and Reality* [Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume XIX]. Reading, Massachusetts: Addison Wesley.

Demetrius. *On Style*. 1995 [1932]. Edited and translated by Doreen C. Innes; based on the translation by W. Rhys Roberts. In *Aristotle*, volume 23. Cambridge: Harvard [Loeb].

Fahnestock, Jeanne. (in press). *Figures of Argument: Studies in the Rhetoric of Science*. New York: Oxford University Press.

Fauconnier, Gilles. *Mappings in Thought and Language.* Cambridge: Cambridge University Press, 1997.

Fauconnier, Gilles and Mark Turner. 1994. "Conceptual projection and middle spaces," UCSD Cognitive Science Technical Report 9401. San Diego. [Available from http://cogsci.ucsd.edu and from http://blending.stanford.edu]

_____. 1996. "Blending as a Central Process of Grammar" in *Conceptual Structure, Discourse, and Language*. Edited by Adele Goldberg. Stanford: Center for the Study of Language and Information.

_____. (1998a). "Principles of Conceptual Integration" in *Discourse and Cognition.* Edited by Jean-Pierre Koenig. Stanford: Center for the Study of Language and Information, 1998, pages 269-283.

_____. 1998b. "Conceptual Integration Networks." *Cognitive Science* 22:2 (April-June 1998), 133-187.

_____. 2002. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities.* New York: Basic Books, 2002.

Fillmore, Charles, Paul Kay, and Catherine O'Connor. 1988. "Regularity and idiomaticity in grammatical constructions: The Case of *Let Alone*.." *Language* 64, 501-538.

Freeman, Margaret. 1997. "Grounded spaces: Deictic -self anaphors in the poetry of Emily Dickinson," *Language and Literature*, 6:1, pages 7-28.

Gagné, Christina L. and Gregory L. Murphy. 1996. "Influence of Discourse Context on Feature Availability in Conceptual Combination." *Discourse Processes* 22, 79-101.

Geertz, Clifford. 1973. *The Interpretation of Cultures*. New York: Basic.

Goldberg, Adele. 1995. *Constructions: A Construction Grammar Approach to Argument Structure.* Chicago: University of Chicago Press.

Goldberg, Adele. (1996) "Construction Grammar." *Concise Encyclopedia of Syntactic Theories.* Brown and Miller, editors. Oxford: Elsevier Science Limited.

Hermogenes. 1913. *Peri Ideon* in *Hermogenis Opera.* Edited by Hugo Rabe. Leipzig: Teubner.

Hermogenes. 1987. *On Types of Style*. Translated by Cecil W. Wooten. Chapel Hill: University of North Carolina Press.

Hogan, Lisa. 1995. "Novel blends in contemporary English," manuscript, University of New Mexico.

Hubel, David H. 1995. *Eye, Brain, and Vision.* New York: Scientific American Library.

Israel, Michael. 1996. "The *Way* Constructions Grow." In *Conceptual Structure, Discourse, and Language*. Edited by Adele Goldberg. Stanford: Center for the Study of Language and Information.

Johnson, Mark. 1987. *The Body in the Mind*. Chicago: University of Chicago Press.

Kay, Paul. 1995. "Construction Grammar." *Handbook of pragmatics*. Edited by Jef Verschueren, Jan-Ola Ostman, and Jan Blommaert. Amsterdam; Philadelphia: J. Benjamins.

Kay, Paul and Charles Fillmore. (manuscript, 1995) "Grammatical Constructions and Linguistic Generalizations: the *What's X doing Y?* Construction." Unpublished manuscript. Department of Linguistics, UC Berkeley.

Kennedy, George, editor, commentator, and translator. (1991) Aristotle. *On rhetoric : a theory of civic discourse* . New York : Oxford University Press.

Lakoff, George. 1987. *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind.* Chicago: University of Chicago Press.

Lakoff, George & Johnson, Mark. 1980. *Metaphors We Live By*. Chicago: University of Chicago Press.

Lewis, C. S. 1936. *The Allegory of Love.* Oxford: Oxford University Press.

Longinus. 1995 [1932]. *On the Sublime.* Translated by W. Hamilton Fyfe and revised by Donald Russell. In *Aristotle*, volume 23. Cambridge: Harvard University Press [Loeb].

Mandelblit, Nili. 1995. "Beyond Lexical Semantics: Mapping and Blending of Conceptual and Linguistic Structures in Machine Translation." In *Proceedings of the Fourth International Conference on the Cognitive Science of Natural Language Processing*, Dublin.

Mandelblit, Nili. 1997. "Grammatical Blending: Creative and Schematic Aspects in Sentence Processing and Translation." Ph.D. dissertation, UC San Diego.

Oakley, Todd. 1995. "Presence: the conceptual basis of rhetorical effect." University of Maryland dissertation.

Quinn, Arthur. 1982. *Figures of Speech: Sixty Ways to Turn a Phrase.* Salt Lake City: Gibbs M. Smith, Inc.

Quintilian. 1921. *Institutio Oratoria.* Four volumes. Translated by H. E. Butler. Cambridge: Harvard University Press [Loeb].

Ramey, Martin. 1997. "Eschatology and Ethics" in "The Problem Of The Body: The Conflict Between Soteriology and Ethics In Paul." Doctoral dissertation, Chicago Theological Seminary.

Richards, I. A. 1936. "Metaphor," a lectured delivered in 1936 and printed as chapter five of *The Philosophy of Rhetoric.* New York: Oxford University Press.

Sun, Douglas. 1994. "Thurber's Fables for our Time: A Case Study in Satirical Use of the Great Chain Metaphor" in *Studies in American Humor*, new series volume 3, number 1, pages 51-61.

Sweetser, Eve. 1990. *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure.* Cambridge: Cambridge University Press.

Tabor, Whitney. 1994a. Syntactic innovation: A connectionist model. Ph.D. dissertation, Stanford University.

Tabor, Whitney. 1994b. "The gradual development of degree modifier *sort of*: A corpus proximity model. In K. Beals, G. Cooke, D. Kathman, K.-E. McCullough, S. Kita, and D. Testen, editors, *Proceedings of the 29th Regional Meeting of the Chicago Linguistic Society.* University of Chicago.

Thomas , Francis-Noël and Mark Turner. 1994. *Clear and Simple as the Truth: Writing Classic Prose.* Princeton: Princeton University Press, 1994.

Todorov, Tzvetan. 1982. *Theories of the Symbol.* Translated by Catherine Porter. Ithaca: Cornell University Press.

Traugott, Elizabeth. 1995. "Subjectification in grammaticalization." In Dieter Stein and Susan Wright, editors, *Subjectivity and Subjectification in Language.* Cambridge: Cambridge University Press. Pages 31-54.

Traugott, Elizabeth. 1996. "On the role of constructions in grammaticalization." Brian Joseph and Richard Janda, editors, *Handbook of historical linguistics*. Oxford: Blackwell.

Turner, Mark. 1996a. "Conceptual Blending and Counterfactual Argument in the Social and Behavioral Sciences," Philip Tetlock and Aaron Belkin, editors, *Counterfactual Thought Experiments in World Politics*. Princeton, N.J.: Princeton University Press.

Turner, Mark. 1996b. *The Literary Mind*. New York: Oxford University Press.

Turner, Mark. 1991. *Reading Minds: The Study of English in the Age of Cognitive Science*. Princeton: Princeton University Press.

Turner, Mark. 1987. *Death is the Mother of Beauty: Mind, Metaphor, Criticism*. Chicago: University of Chicago Press.

Turner, Mark. 1989. "Categories and Analogies" in *Analogical Reasoning: Perspectives of Artificial Intelligence, Cognitive Science, and Philosophy*. Edited by David Helman. Dordrecht: Kluwer, 3-24.

Turner, Mark and Gilles Fauconnier. 1995. "Conceptual Integration and Formal Expression." *Metaphor and Symbolic Activity*. 10:3, 183-203.

_____. 1998 "Conceptual Integration in Counterfactuals" in *Discourse and Cognition.* Edited by Jean-Pierre Koenig. Stanford: Center for the Study of Language and Information, pages 285-296.

Veale, Tony. 1996. "Pastiche: A Metaphor-centred Computational Model of Conceptual Blending, with special reference to Cinematic Borrowing." Manuscript. Available from the Conceptual Integration Web Page: http://blending.stanford.edu

Vickers, Brian. 1988. *In Defence of Rhetoric.* Oxford: Clarendon Press.

Zbikowski, Lawrence. 1996. "Conceptual blending and song." Department of Music, University of Chicago.

Zeki, Semir. 1993. *A Vision of the Brain.* Oxford: Blackwell.