

Conceptual Integration Networks

[Expanded web version, 10 February 2001]

Gilles Fauconnier & Mark Turner

The web page for research on conceptual integration is <http://blending.stanford.edu>

Original version published in *Cognitive Science*, 22(2) 1998, 133-187.

Copyright © Cognitive Science Society, Inc. Used by permission.

Abstract

Conceptual integration—"blending"—is a general cognitive operation on a par with analogy, recursion, mental modeling, conceptual categorization, and framing. It serves a variety of cognitive purposes. It is dynamic, supple, and active in the moment of thinking. It yields products that frequently become entrenched in conceptual structure and grammar, and it often performs new work on its previously entrenched products as inputs. Blending is easy to detect in spectacular cases but it is for the most part a routine, workaday process that escapes detection except on technical analysis. It is not reserved for special purposes, and is not costly.

In blending, structure from input mental spaces is projected to a separate, "blended" mental space. The projection is selective. Through completion and elaboration, the blend develops structure not provided by the inputs. Inferences, arguments, and ideas developed in the blend can have effect in cognition, leading us to modify the initial inputs and to change our view of the corresponding situations.

Blending operates according to a set of uniform structural and dynamic principles. It additionally observes a set of optimality principles.

Contents

- I. Introduction
- II. An illustration
- III. The network model of conceptual integration
- IV. Applications
- V. Advanced aspects of the network model
- VI. Optimality principles
- VII. Additional dimensions of conceptual integration
- VIII. Summary and further results
- IX. Conclusion

I. Introduction

Much of the excitement about recent work on language, thought, and action stems from the discovery that the same structural cognitive principles are operating in areas that were once viewed as sharply distinct and technically incommensurable. Under the old view, there were word meanings, syntactic structures, sentence meanings (typically truth-conditional), discourse and pragmatic principles, and then, at a higher level, figures of speech like metaphor and metonymy, scripts and scenarios, rhetoric, forms of inductive and deductive reasoning, argumentation, narrative structure, etc. A recurrent finding in recent work has been that key notions, principles, and instruments of analysis cut across all these divisions and in fact operate in non-linguistic situations as well. Here are some of them:

Frames structure our conceptual and social life. As shown in the work of Fillmore, Langacker, Goldberg, and others, they are also, in their most generic, and schematic forms, a basis for grammatical constructions. Words are themselves viewed as constructions, and lexical meaning is an intricate web of connected frames. Furthermore, although cognitive framing is reflected and guided by language, it is not inherently linguistic. People manipulate many more frames than they have words and constructions for.

Analogical mapping, traditionally studied in connection with reasoning, shows up at all levels of grammar and meaning construction, such as the interpretation of counterfactuals and hypotheticals, category formation, and of course metaphor, whether creative or conventional.

Reference points, focus, viewpoints, and dominions are key notions not only at higher levels of narrative structure, but also at the seemingly micro-level of ordinary grammar, as shown convincingly by Langacker 1993, Zribi-Hertz 1989, Van Hoek 1997, Cutrer 1994, among others.

Connected mental spaces account for reference and inference phenomena across wide stretches of discourse, but also for sentence-internal multiple readings and tense/mood distributions. Mappings at all levels operate between such spaces, and like frames they are not specifically linguistic. (Fauconnier 1997, Dinsmore 1991, Cutrer 1994, Fauconnier and Sweetser, 1996).

Connectors and conceptual connections also operate at all levels, linking mental spaces and other domains for coreference, for metonymy (Nunberg 1978), and for analogy and metaphor (Turner 1991, Sweetser 1990).

There are other notions that apply uniformly at seemingly different levels, such as figure/ground organization (Talmy 1978), profiling, or pragmatic scales. Running through this research is the central cognitive scientific idea of *projection* between structures. Projection connects frames to specific situations, to related frames, and to conventional scenes. Projection connects related linguistic constructions. It connects one viewpoint to another and sets up new viewpoints partly on the basis of old. It connects counterfactual conceptions to

non-counterfactual conceptions on which they are based. Projection is the backbone of analogy, categorization, and grammar.

In the present study, we show that *projection typically involves conceptual integration*. There is extensive previous research on varieties of projection, but not on conceptual integration. Empirical evidence suggests that an adequate characterization of mental projection requires a theory of conceptual integration. We propose the basis for such a theory and argue that conceptual integration—like framing or categorization—is a basic cognitive operation that operates uniformly at different levels of abstraction and under superficially divergent contextual circumstances. It also operates along a number of interacting gradients. Conceptual integration plays a significant role in many areas of cognition. It has uniform, systematic properties of structure and dynamics.

The nature of mapping between domains has enjoyed sustained attention as a central problem of cognitive science, and voluminous literatures have developed in this area, including studies by those who call their subject "analogy" or "similarity" (e. g., Hofstadter 1985, 1995a, Mitchell 1993, French 1995, Keane, Ledgeway, and Duff 1994; Holyoak and Thagard, 1989, 1984; Forbus, Gentner, and Law, 1994; Gentner 1983, 1989; Holland, Holyoak, Nesbett, and Thagard, 1986), studies by those who call their subject "metaphor" (e.g., Lakoff and Johnson 1980; Lakoff and Turner 1989; Sweetser 1990; Turner 1987; Indurkha 1992; Gibbs 1994) and studies that consider cross-domain mapping in general (e.g., Fauconnier 1997, Ortony 1979a, 1979b, Glucksberg and Keysar 1990, Turner 1991).

Our immediate goal is *not* to take a stand on issues and problems of cross-space mappings. Those issues are many and the debates over them will continue and will be further enriched, we hope, by taking blending into consideration. What we *will* be suggesting is that models of cross-space mapping do not by themselves explain the relevant data. These data involve conceptual integration and multiple projections in ways that have typically gone unnoticed. Cross-space mapping is only one aspect of conceptual integration, and the existing body of research on the subject overlooks conceptual integration, which it is our intention to foreground and analyze here. As we move through the data that crucially involves both cross-space mapping and conceptual integration, we will remark that much of it is neither metaphoric nor analogical. [\[1\]](#)

We take it as an established and fundamental finding of cognitive science that structure mapping and metaphorical projection play a central role in the construction of reasoning and meaning. In fact, the data we analyze shows that such projections are even more pervasive than previously envisioned. Given the existence and key role of such mappings, our focus is on the construction of additional spaces with emergent structure, not directly available from the input domains.

We also rely on another fundamental finding of cognitive science, the capacity for mental simulation, as demonstrated in Johnson-Laird (1983), Kahneman (1995), Grush (1995), Schwartz and Black (1996), Barsalou (1996) among others. In our analysis, the simulation capacity assists in the on-line elaboration of blended spaces ("running the blend"). There is the added twist that simulation can operate on mental spaces which need not have potential real world reference.

Our methodology and argumentation take the following form. Since the cognitive process of conceptual integration has been largely overlooked, it is useful to give evidence for its operation in a wide variety of areas. Since conceptual integration has uniform structural and dynamic properties, it is important to reveal this uniformity behind the appearance of observational and functional diversity. We proceed analytically and empirically, by showing that central inferences, emotions, and conceptualizations, not explained in currently available frameworks, are accounted for elegantly by the conceptual integration model. The argumentation often takes the following specific form: a particular process of meaning construction has particular input representations; during the process, inferences, emotions and event-integrations emerge which cannot reside in any of the inputs; they have been constructed dynamically in a new mental space—the blended space—linked to the inputs in systematic ways. For example, "They dug their own financial grave" draws selectively from different and incompatible input frames to construct a blended space that has its own emergent structure and that provides central inferences. In this case, the blended space has become conventional.

The diversity of our data (of which only a small sample appears in the present paper) is necessary to support our claim for generality. (In showing that cell division is a basic process, it is necessary to study it for many kinds of cells. In arguing that natural selection is a general principle, it is necessary to exemplify it for widely different organisms and species.) In arguing that conceptual integration is a basic cognitive operation, we must show that it operates in many different kinds of cases.

Conceptual blending is not a compositional algorithmic process and cannot be modeled as such for even the most rudimentary cases. Blends are not predictable solely from the structure of the inputs. Rather, they are highly motivated by such structure, in harmony with independently available background and contextual structure; they comply with competing optimality constraints discussed in section VI, and with locally relevant functional goals. In this regard, the most suitable analog for conceptual integration is not chemical composition but biological evolution. Like analogy, metaphor, translation, and other high-level processes of meaning construction, integration offers a formidable challenge for explicit computational modeling.

Special cases of conceptual blending have been discussed insightfully by Koestler (1964), Goffman (1974), Talmy (1977), Fong (1988), Moser and Hofstadter (ms.), and Kunda, Miller and Clare (1990). Fauconnier (1990) and Turner (1991) also contain analyses of such phenomena. All these authors, however, take blends to be somewhat exotic, marginal manifestations of meaning. We will show here that the process is in fact central, uniform, and pervasive.

The data and analysis we consider here suggest many psychological and neuropsychological experiments (Coulson 1997), but in the present work our emphasis is on the understanding of ecologically valid data. Research on meaning, we suggest, requires analysis of extensive ranges of data, which must be connected theoretically across fields and disciplines by general cognitive principles.

We start our report with an effective but somewhat idealized example of blending, in order to illustrate the issues and terminology. We then outline the general process of conceptual integration and the systematic dynamic properties of blends. We work through some case-studies in a variety of areas. Section VI presents

the competing optimality principles under which conceptual integration operates.

II. An illustration

The riddle of the Buddhist monk

Consider a classic puzzle of inferential problem-solving (Koestler 1964):

A Buddhist monk begins at dawn one day walking up a mountain, reaches the top at sunset, meditates at the top for several days until one dawn when he begins to walk back to the foot of the mountain, which he reaches at sunset. Making no assumptions about his starting or stopping or about his pace during the trips, prove that there is a place on the path which he occupies at the same hour of the day on the two separate journeys.

Our demonstration of the power of blending is likely to be more effective if the reader will pause for a moment and try to solve the problem before reading further. The basic inferential step to showing that there is indeed such a place, occupied at exactly the same time going up and going down, is to imagine the Buddhist monk walking both up and down the path on the same day. Then there must be a place where he meets himself, and that place is clearly the one he would occupy at the same time of day on the two separate journeys.

The riddle is solved, but there is a cognitive puzzle here. The situation that we devised to make the solution transparent is a fantastic one. The monk cannot be making the two journeys simultaneously on the same day, and he cannot "meet himself." And yet this implausibility does not stand in the way of understanding the riddle and its solution. It is clearly disregarded. The situation imagined to solve the riddle is a blend: it combines features of the journey to the summit and of the journey back down, and uses emergent structure in that blend to make the affirmative answer apparent. Here is how this works.

Mental space. In our model, the input structures, generic structures, and blend structures in the network are *mental spaces*. Mental spaces are small conceptual packets constructed as we think and talk, for purposes of local understanding and action. Mental spaces are very partial assemblies containing elements, and structured by frames and cognitive models. They are interconnected, and can be modified as thought and discourse unfold. Mental spaces can be used generally to model dynamical mappings in thought and language. Fauconnier (1994), Fauconnier (1997), Fauconnier & Sweetser (1996).

In our diagrams, the mental spaces are represented by circles, elements by points (or sometimes icons) in the circles, and connections between elements in different spaces by lines. The frame structure recruited to the mental space is represented either outside in a rectangle or iconically inside the circle.

Input spaces. There are two input spaces. Each is a partial structure corresponding to one of the two journeys.




Figure 1

d₁ is the day of the upward journey, and d₂ the day of the downward journey. a₁ is the monk going up, a₂ is the monk going down.

Cross-space mapping of counterpart connections. There is a partial cross-space mapping between the input spaces. The cross-space mapping connects counterparts in the input spaces. It connects the mountain, moving individual, day of travel, and motion in one space to the mountain, moving individual, day, and motion in the other space.




Figure 2

Generic space. There is a generic space, which maps onto each of the inputs. The generic space contains what the inputs have in common: a moving individual and his position, a path linking foot and summit of the mountain, a day of travel. It does not specify the direction of motion or the actual day. (At this point in our exposition, it will not be clear why our model needs a generic space in addition to a cross-space mapping. Later, we will argue that powerful generic spaces can themselves become conventional and serve as resources to be drawn on in attempts to build new cross-space mappings in new integration networks.)

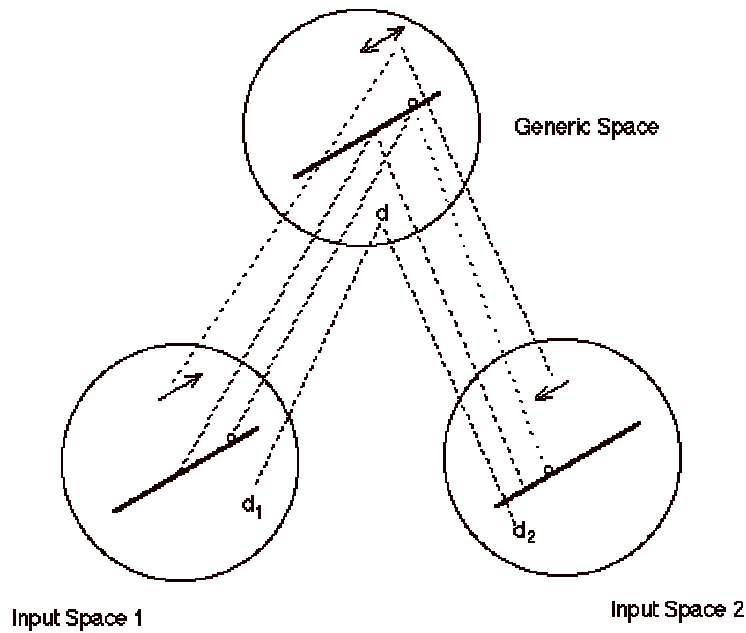


Figure 3

Blend. There is a fourth space, the blend. In the blend, the two counterpart identical mountain slopes are mapped onto a single slope. The two days of travel, d_1 and d_2 , are mapped onto a single day d' and therefore fused. While in the generic space and each of the input spaces there is only one moving individual, in the blend there are two moving individuals. The moving individuals in the blend and their positions have been projected from the inputs in such a way as to preserve time of day and direction of motion, and *therefore the two moving individuals cannot be fused*. Input 1 represents dynamically the entire upward journey, while Input 2 represents the entire downward journey. The projection into the blend preserves times and positions. The blend at time t of day d' contains a counterpart of a_1 at the position occupied by a_1 at time t of d_1 , and a counterpart of a_2 at the position occupied by a_2 at time t of day d_2 .

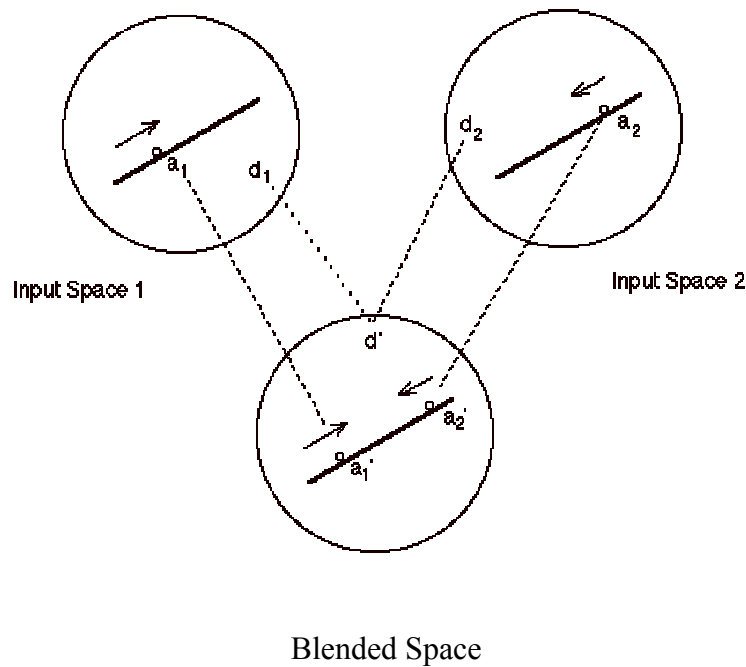


Figure 4

Selective projection. The projection of structure to the blend is selective. For example, the calendrical time of the journey is not projected to the blend.

Emergent structure. The blend contains emergent structure not in the inputs. First, *composition* of elements from the inputs makes relations available in the blend that did not exist in the separate inputs. In the blend but in neither of the inputs, there are two moving individuals instead of one. They are moving in opposite directions, starting from opposite ends of the path, and their positions can be compared at any time of the trip, since they are traveling on the same day, d' .

Second, *completion* brings additional structure to the blend. This structure of two people moving on the path can itself be viewed as a salient part of a familiar background frame: two people starting a journey at the same time from opposite ends of a path. By *completion*, this familiar structure is recruited into the blend. We know, from "common sense," i.e. familiarity with this background frame, that the two people will necessarily meet at some time t' of their journey. We do not have to compute this encounter afresh; it is supplied by completion from a pre-existing familiar frame. There is no encounter in the generic space or either of the inputs, but there is an encounter in the blend, and it supplies the central inference.

Importantly, the blend remains hooked up to the Inputs, so that structural properties of the blend can be mapped back onto the Inputs. In our example, because of the familiarity of the frame obtained by completion, the inference that there is a meeting time t' with a common position p is completely automatic. The mapping back to the input spaces yields:

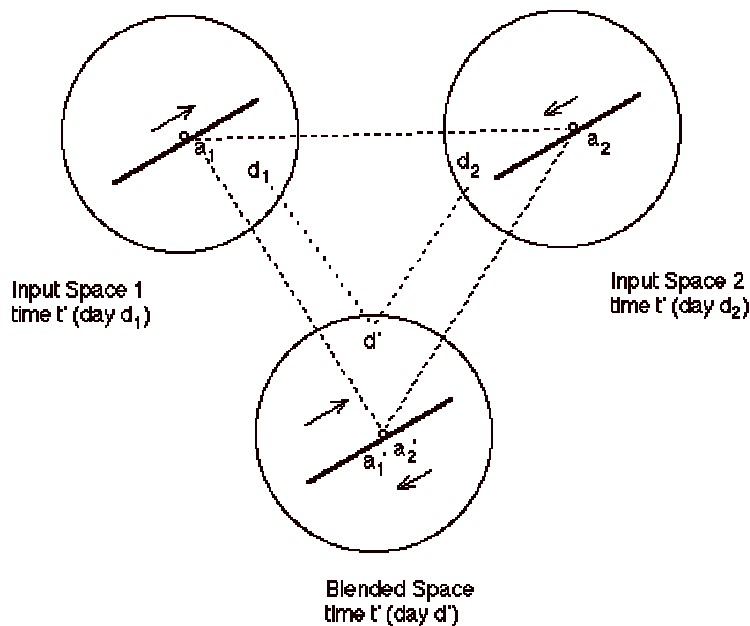


Figure 5

Since the projection of individuals into the blend preserves positions on the path, we "know" through this mapping that the positions of a_1 and a_2 are the "same" at time t' on the different days, simply because they are the same, by definition, in the frame of two people meeting, instantiated in the blend by their counterparts a_1' and a_2' .

It is worth emphasizing that the pragmatic incongruity in the blend of the same person traveling in two opposite directions and meeting himself is disregarded, because the focus of the problem is the meeting point and its counterparts in the Input spaces. Blends are used cognitively in flexible ways. By contrast, in examples we discuss later, similar incongruities in the blend get highlighted and mapped back to the Inputs for inferential and emotional effect. Incongruity makes blends more visible, but blends need not be incongruous—incongruity is not one of their defining characteristics.

Notice also that in this blend, some counterparts have been fused (the days, the path on the different days, and the corresponding times on different days), others have been projected separately (the monk on the way up, the monk on the way down, the directions of motion). Projection from the Inputs is only partial—the specific dates of the journeys are not projected, nor the fact that the monk will stay at the top for a while after his upward journey. But the blend has new "emergent" structure not in the Inputs: two moving individuals whose positions can be compared and may coincide, and the richer frame of two travelers going in opposite directions on the same path and necessarily meeting each other. This emergent structure is crucial to the performance of the reasoning task.

Rather amazingly, the Buddhist monk blend shows up in real life. Hutchins (1995) studies the fascinating mental models set up by Micronesian navigators to sail across the Pacific. In such models, it is the islands that move, and virtual islands may serve as reference points. Hutchins reports a conversation between Micronesian and Western navigators who have trouble understanding each other's conceptualizations. As described in Lewis (1972), the Micronesian navigator Beiong succeeds in understanding a Western diagram of intersecting bearings in the following way:

He eventually succeeded in achieving the mental tour de force of visualizing himself sailing simultaneously from Oroluk to Ponape and from Ponape to Oroluk and picturing the ETAK bearings to Ngatik at the start of both voyages. In this way he managed to comprehend the diagram and confirmed that it showed the island's position correctly. [the etak is the virtual island, and Ngatik is the island to be located.]

Previous insightful work by Kahneman (1995), Schwartz and Black (1996), Barsalou (1996), has emphasized the role of imaginative mental simulation and depiction in making inferences about physical scenarios. In the riddle of the Buddhist Monk, the physical system we are interested in consists of the sequence of the monk's departing, traveling up the hill, reaching the top, waiting, departing, traveling down the hill, and reaching the bottom. Imagining a mental depiction of this scenario does not solve the riddle, but representing it isomorphically as two input spaces to a blend and imagining a mental depiction of that blend does indeed create an event of encounter in the blend which points to a solution, not for the blend, but for the input spaces and therefore identically for the original scenario. Mental simulation, in this case, depends indispensably

upon conceptual blending to provide the effective scenario to begin with.

III. The network model of conceptual integration

In this section, we present the central features of our network model, keyed to the illustration we have just given. In section V, we present advanced aspects of the model.

The network model is concerned with on-line, dynamical cognitive work people do to construct meaning for local purposes of thought and action. It focuses specifically on conceptual projection as an instrument of on-line work. Its central process is conceptual blending.




Figure 6

Mental spaces. The circles in Figure 6 represent mental spaces. In the monk example, there are four mental spaces: the two inputs, the generic, and the blend. There are also background frames recruited to build these mental spaces, such as the background frame of two people approaching each other on a path. This is a minimal network. Networks in other cases of conceptual integration may have yet more input spaces and even multiple blended spaces.

Cross-space mapping of counterpart connections. In conceptual integration, there are partial counterpart connections between input spaces. The solid lines in Figure 6 represent counterpart connections. Such counterpart connections are of many kinds: connections between frames and roles in frames; connections of identity or transformation or representation; metaphoric connections, etc. In the monk example, the monks, paths, journeys, days, and so on are counterparts.

Generic space. As conceptual projection unfolds, whatever structure is recognized as belonging to both of the input spaces constitutes a generic space. At any moment in the construction, the generic space maps onto each of the inputs. It defines the current cross-space mapping between them. A given element in the generic space maps onto paired counterparts in the two input spaces.

Blending. In blending, structure from two input mental spaces is projected to a third space, the "blend." In the monk example, the two input spaces have two journeys completely separated in time; the blend has two simultaneous journeys. Generic spaces and blended spaces are related: blends contain generic structure captured in the generic space, but also contain more specific structure, and can contain structure that is impossible for the inputs, such as two monks who are the same monk.

Selective projection. The projection from the inputs to the blend is typically partial. In Figure 6, not all elements from the inputs are projected to the blend.

There are three operations involved in constructing the blend: composition, completion, and elaboration.

Composition. Blending composes elements from the input spaces, providing relations that do not exist in the separate inputs. In the monk riddle, composition yields two travelers making two journeys. Fusion is one kind of composition. Counterparts may be brought into the blend as separate elements or as a fused element. Figure 6 represents one case in which counterparts are fused in the blend and one case in which counterparts are brought into the blend as distinct entities. In the monk example, the two days in the inputs are fused into one day in the blend, but the two monks from the inputs are brought into the blend as distinct entities.

Completion. Blends recruit a great range of background conceptual structure and knowledge without our recognizing it consciously. In this way, composed structure is completed with other structure. The fundamental subtype of recruitment is pattern completion. A minimal composition in the blend can be extensively completed by a larger conventional pattern. In the monk example, the structure achieved through

composition is completed by the scenario of two people journeying toward each other on a path, which yields an encounter.

Elaboration. Elaboration develops the blend through imaginative mental simulation according to principles and logic in the blend. Some of these principles will have been brought to the blend by completion. Continued dynamic completion can recruit new principles and logic during elaboration. But new principles and logic may also arise through elaboration itself. We can "run the blend" indefinitely: for example, the monks might meet each other and have a philosophical discussion about the concept of identity. Blended spaces can become extremely elaborated.

Emergent structure. Composition, completion, and elaboration lead to emergent structure in the blend; the blend contains structure that is not copied from the inputs. In Figure 6, the square inside the blend represents emergent structure.

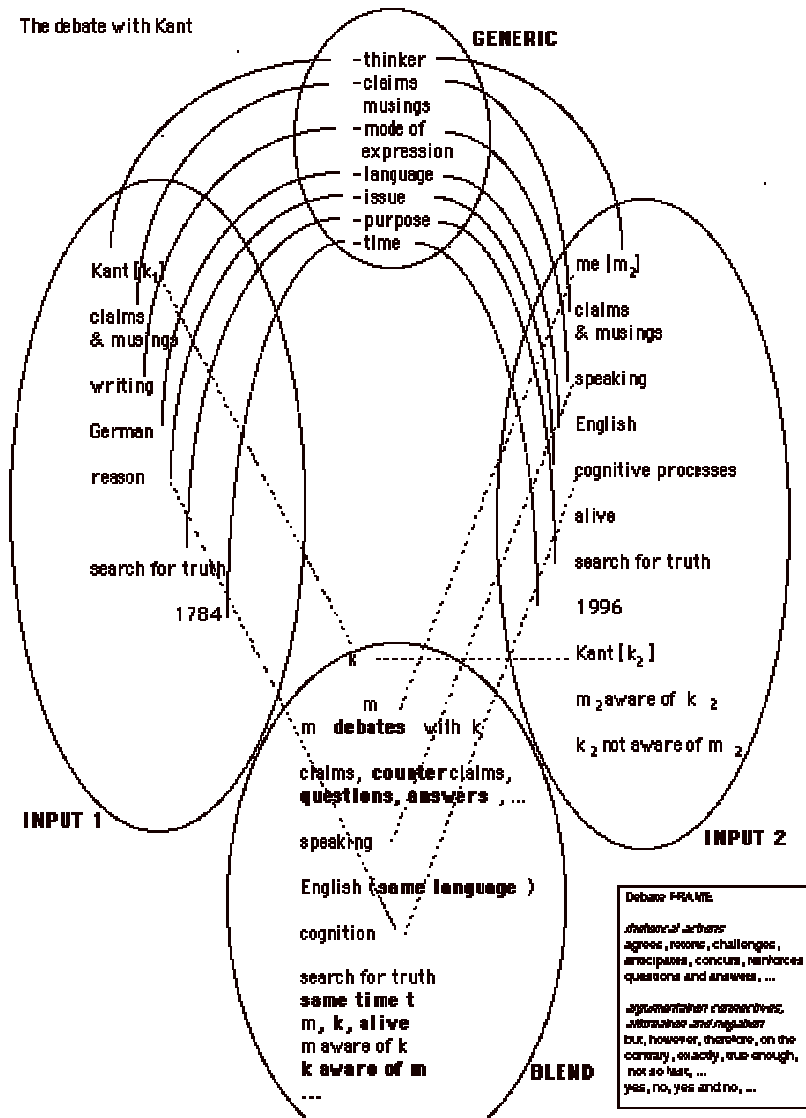
IV. Applications

The Debate with Kant

The monk example presents a salient and intuitively apparent blend, precisely because of its pragmatic anomaly. But our claim is that blends abound in all kinds of cases that go largely unnoticed. Some are created as we talk, others are conventional, and others are even more firmly entrenched in the grammatical structure. We discuss in Fauconnier & Turner 1996, the situation in which a contemporary philosopher says, while leading a seminar,

I claim that reason is a self-developing capacity. Kant disagrees with me on this point. He says it's innate, but I answer that that's begging the question, to which he counters, in *Critique of Pure Reason*, that only innate ideas have power. But I say to that, what about neuronal group selection? He gives no answer.

In one input mental space, we have the modern philosopher, making claims. In a separate but related input mental space, we have Kant, thinking and writing. In neither input space is there a debate. The blended space has both the modern philosopher (from the first input space) and Kant (from the second input space). In the blend, the additional frame of *debate* has been recruited, to frame Kant and the modern philosopher as engaged in simultaneous debate, mutually aware, using a single language to treat a recognized topic.



The debate frame comes up easily in the blend, through pattern completion, since so much of its structure is already in place in the composition of the two inputs. Once the blend is established, we can operate cognitively within that space, which allows us to manipulate the various events as an integrated unit. The debate frame brings with it conventional expressions, available for our use. We know the connection of the blend to the input spaces, and the way in which structure or inferences developed in the blend translate back to the input spaces.

A "realist" interpretation of the passage would be quite fantastic. The philosophy professor and Kant would have to be brought together in time, would have to speak the same language, and so on. No one is fooled into

thinking that this is the intended interpretation. In fact, using a debate blend of this type is so conventional that it will go unnoticed.

And yet, it has all the criterial properties of blending. There is a Cross-space mapping linking Kant and his writings to the philosophy professor and his lecture. Counterparts include: Kant and the professor, their respective languages, topics, claims, times of activity, goals (e.g. search for truth), modes of expression (writing vs. speaking).

There is Partial projection to the blend: Kant, the professor, some of their ideas, and the search for truth are projected to the blend. Kant's time, language, mode of expression, the fact that he's dead, and the fact that he was never aware of the future existence of our professor are not projected.

There is Emergent Structure through Composition: we have two people talking in the same place at the same time. There is Emergent Structure through Completion: two people talking in the same place at the same time evoke the cultural frame of a conversation, a debate (if they are philosophers), an argument. This frame, the debate frame, structures the blend and is reflected by the syntax and vocabulary of the professor (*disagrees, answer, counters, what about, ...*).

This example allows us to observe that blends provide Integration of Events: Kant's ideas and the professor's claims are integrated into a unified event, the debate. Looking back now to the monk example, we see that the blend in that case integrated into a single scenario various events of uncertain relation spread out over time. Blends provide a space in which ranges of structure can be manipulated uniformly. The other spaces do not disappear once the blend has been formed. On the contrary, the blend is valuable only because it is connected conceptually to the inputs. The monk blend tells us something about the inputs. The debate with Kant tells us something about the inputs.

Complex numbers

Conceptual projection enables us to extend categories to cover new provisional members. The blended space that develops during such a projection merges the original category with its new extension. When categories are extended permanently, it is the structure of this blend that defines the new category structure, thus carving out a novel conceptual domain. The history of science, and of mathematics and physics in particular, is rich in such conceptual shifts. (See Fauconnier & Turner 1994; Lakoff & Núñez in preparation; Lansing, personal communication.) It is customary to speak of models either replacing or extending previous models, but the pervasiveness and importance of merging may have been underestimated.

Consider as an example the stage of mathematical conceptual development at which complex numbers became endowed with angles (arguments) and magnitudes. Square roots of negative numbers had shown up in formulas of sixteenth-century mathematicians and operations on these numbers had been correctly

formulated. But the very mathematicians who formulated such operations, Cardan and especially Bombelli, were also of the opinion that they were "useless," "sophistic," and "impossible" or "imaginary." Such was also the opinion of Descartes a century later. Leibniz said no harm came of using them, and Euler thought them impossible but nevertheless useful. The square roots of negative numbers had the strange property of lending themselves to formal manipulations without fitting into a mathematical conceptual system. A genuine concept of complex number took time to develop, and the development proceeded in several steps along the lines explained above for analogical connections and blending.

The first step exploited the preexisting analogical mapping from numbers to one-dimensional space. Wallis is credited with having observed—in his *Algebra* (1685)—that if negative numbers could be mapped onto a directed line, complex numbers could be mapped onto points in a two-dimensional plane, and he provided geometrical constructions for the counterparts of the real or complex roots of $ax^2 + bx + c = 0$ (Kline 1980). In effect, Wallis provided a model for the mysterious numbers, thereby showing their consistency, and giving some substance to their formal manipulation. This is of course a standard case of extending analogical connections; geometric space is a source domain partially mapped onto the target domain of numbers. The mapping from a single axis is extended to mapping from the whole plane; some geometric constructions are mapped onto operations on numbers. Notice that neither the original mapping nor its extension requires more than two domains. We do not need a generic space, since there is no assumption in work like Wallis's that numbers and points in a plane share properties at some higher level of abstraction. The necessary structure is already present in the conceptual domain of two-dimensional space because it already contains the notion of distance which is expressed directly by means of numbers. (Of course, this source domain has a conceptual history of its own. We argue elsewhere that in fact it is itself the product of a non-trivial conceptual blend.) Nor does it involve a blend; numbers and points remain totally distinct categories at all levels. Although the mapping proposed by Wallis showed the formal consistency of a system including complex numbers, it did not provide a new extended concept of number. As Morris Kline reports, Wallis's work was ignored: it did not make mathematicians receptive to the use of such numbers. In itself, this is an interesting point. It shows that mapping a coherent space onto a conceptually incoherent space is not enough to give the incoherent space new conceptual structure. It also follows that coherent abstract structure is not enough, even in mathematics, to produce satisfactory conceptual structure: In Wallis's representation, the metric geometry provided abstract schemas for a unified interpretation of real and imaginary numbers, but this was insufficient cognitively for mathematicians to revise their domain of numbers accordingly.

In the analysis developed here, the novel conceptual structure in the mathematical case of numbers is first established within a blended space. In the blend, but not in the original inputs, it is possible for an element to be simultaneously a number and a geometric point, with cartesian coordinates (a,b) and polar coordinates (r,φ) . In the blend, we find interesting general formal properties of such numbers, such as

$$(a, b) + (a', b') = (a+a', b+b')$$

$$(r, \varphi) \times (r', \varphi') = (rr', \varphi + \varphi')$$

Every number in this extended sense has a real part, an imaginary one, an argument, and a magnitude. By virtue of the link of the blend to the geometric input space, the numbers can be manipulated geometrically; by virtue of the link of the blend to the input space of real numbers, the new numbers in the blend are immediately conceptualized as an extension of the old numbers (which they include by way of the mapping). As in Wallis's scheme, the mapping from points on a line to numbers has been extended to a mapping from points in a plane to numbers. This mapping is partial from one input to the other—only one line of the plane is mapped onto the numbers of the target domain—but it is total from the geometric input to the blend: all the points of the plane have counterpart complex numbers. And this in turn allows the blend to incorporate the full geometric structure of the geometric input space.

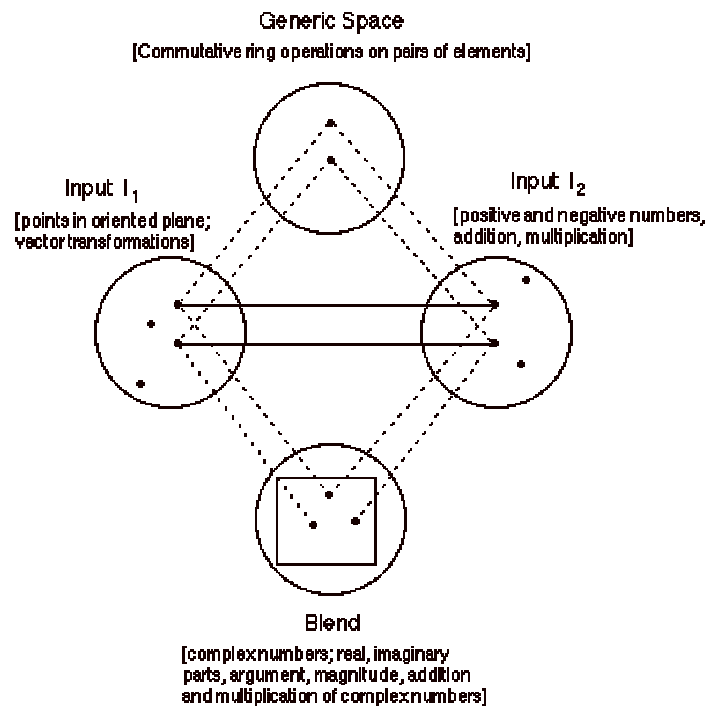


Figure 7

Interestingly, when a rich blended space of this sort is built, an abstract generic space will come along with it. Having the three spaces containing respectively points (input 1), numbers (input 2), complex point/numbers (blend) entails a fourth space with abstract elements having the properties "common" to points and numbers. The relevant abstract notions in this case are those of "operations" on pairs of elements. For numbers, the specific operations (in the target domain) are addition and multiplication. For points in the plane, the operations can be viewed as vector transformations—vector addition, and vector composition by adding angles and multiplying magnitudes. In the blended space of complex numbers, vector addition and

number addition are the same operation, because they invariably yield the same result; similarly, vector transformation and number multiplication are conceptually one single operation. But such an operation can be instantiated algorithmically in different ways depending on which geometric and algebraic properties of the blend are exploited. [2]

In the generic space, specific geometric or number properties are absent. All that is left is the more abstract notion of two operations on pairs of elements, such that each operation is associative, commutative, and has an identity element; each element has under each operation an inverse element; and one of the two operations is distributive with respect to the other. Something with this structure is called by mathematicians a "commutative ring."

The emergence of the concept of complex numbers with arguments and magnitudes displays all the criterial properties of blending. There is an initial cross-space mapping of numbers to geometric space, a generic space, a projection of both inputs to the blend, with numbers fused with geometric points, emergent structure by completion (arguments and magnitudes), and by elaboration (multiplication and addition reconstrued as operations on vectors).

The blend takes on a realist interpretation within mathematics. It constitutes a new and richer way to understand numbers and space. However, it also retains its connections to the earlier conceptions provided by the Input spaces. Conceptual change of this sort is not just replacement. It is the creation of more elaborate and richly connected networks of spaces.

Under our account, then, the evolution and extension of the concept of number includes a four-space stage at which the concept of complex number is logically and coherently constructed in a blended space, on the basis of a generic space structured as a commutative ring. (That generic space is not consciously conceptualized as an abstract domain when the full-blown concept of complex number gets formed. It becomes a conceptual domain in its own right when mathematicians later study it and name it.) The abstract and mathematical example of complex numbers supports the functioning of many-space conceptual projection, with its blended and generic spaces, [3] and confirms that we are dealing with an aspect of thought that is not purely linguistic or verbal. It highlights the deep difference between naming and conceptualizing; adding expressions like $\div -1$ to the domain of numbers, and calling them numbers, is not enough to make them numbers conceptually, even when they fit a consistent model. This is true of category extension in general.

Digging your own grave

Coulson (1997) examines remarkable elaborations of the metaphor "to dig one's own grave." Consider the familiar idiomatic version of the metaphor. "You are digging your own grave" is a conventional expression typically used as a warning or judgment, typically implying that (1) you are doing bad things that will cause you to have a very bad experience, and (2) you are unaware of this causal relation. A conservative parent who keeps his money in his mattress may express

disapproval of an adult child's investing in the stock market by saying, "you are digging your own grave."

At first glance, what we have here is a straightforward projection from the concrete domain of graves, corpses, and burial to abstract domains of getting into trouble, unwittingly doing the wrong things, and ultimate failure. Failing is being dead and buried; bad moves that precede and cause failure are like events (grave-digging) that precede burial. It is foolish to facilitate one's own burial or one's own failure. And it is foolish not to be aware of one's own actions, especially when they are actions leading to one's very extinction.

But a closer look reveals extraordinary mismatches between the purported source and target of this metaphor. The causal structure is inverted. Foolish actions cause failure, but grave-digging does not cause death. It is typically someone's dying that prompts others to dig a grave. And if the grave is atypically prepared in advance, to secure a plot, to keep workers busy, or because the person is expected to die, there is still not the slightest causal connection from the digging to the dying. In the exceptional scenario in which a prisoner is threatened into digging his own grave, it is not the digging that causes the death, and the prisoner will be killed anyway if he refuses. The intentional structure does not carry over. Sextons do not dig graves in their sleep, unaware of what they are doing. In contrast, figurative digging of one's own grave is conceived as unintentional misconstrual of action. The frame structure of agents, patients, and sequence of events is not preserved. Our background knowledge is that the "patient" dies, and then the "agent" digs the grave and buries the "patient." But in the metaphor, the actors are fused and the ordering of events is reversed. The "patient" does the digging, and if the grave is deep enough, has no other option than to die and occupy it. Even in the unusual real life case in which one might dig one's own grave in advance, there would be no necessary temporal connection between finishing the digging and perishing. The internal event structure does not match. In the target, it is certainly true that the more trouble you are in, the more you risk failure. Amount of trouble is mapped onto depth of grave. But again, in the source there is no correlation between the depth of a person's grave and their chances of dying.

Now recall the rationale often proposed for metaphor: Readily available background or experiential structure and inferences of the source are recruited to understand the target. By that standard, and in view of the considerable mismatch, *digging one's own grave* should be a doomed metaphor. In fact, it's a very successful one.

This paradox dissolves when we consider, in addition to the two input spaces, the blended space. In metaphoric cases, such as this one, the two inputs are the "source Input" and the "target Input." The blend in *digging one's own grave* inherits the concrete structure of graves, digging, and burial, from the source Input. But it inherits causal, intentional, and internal event structure from the target Input. They are not simply juxtaposed. Rather, emergent structure specific to the blend is created. In the blend, all the curious properties noted above actually hold. The existence of a satisfactory grave causes death, and is a necessary precondition for it. It follows straightforwardly that the deeper the grave, the closer it is to completion, and the greater the chance for the grave's intended occupant to die. It follows that in the blend (as opposed to the Input source), digging

one's grave is a grave mistake, since it makes dying more probable. In the blend, it becomes possible to be unaware of one's very concrete actions. This is projected from the target Input, where it is indeed fully possible, and frequent, to be unaware of the significance of one's actions. But in the blend, it remains highly foolish to be unaware of such concrete actions; this is projected from the source Input. And it will project back to the target Input to produce suitable inferences (i.e. highlight foolishness and misperception of individual's behavior).

We wish to emphasize that in the construction of the blend, a single shift in causal structure, *the existence of a grave causes death*, instead of *death causes the existence of a grave*, is enough to produce emergent structure, specific to the blend: undesirability of digging one's grave, exceptional foolishness in not being aware of it, correlation of depth of grave with probability of death. The causal inversion is guided by the target, but the emergent structure is deducible within the blend from the new causal structure and familiar common-sense background knowledge. This point is essential, because the emergent structure, although "fantastic" from a literal interpretation point of view, is supremely efficient for the purpose of transferring the intended inferences back to the target Input, and thereby making real-world inferences. This emergent structure is not in the Inputs—it is part of the cognitive construction in the blend. But, also, it is not stated explicitly as part of the blend. It just follows, fairly automatically, from the unstated understanding that the causal structure has been projected from the target, not from the source.

The integration of events in the blend is indexed to events in both of the input spaces. We know how to translate structure in the blend back to structure in the inputs. The blend is an integrated platform for organizing and developing those other spaces. Consider a slightly fuller expression, "with each investment you make, you are digging your grave a little deeper." In the target Input, there are no graves, but there are investments; in the source Input, the graves are not financial, but one does dig; in the blend, investments are simultaneously instruments of digging, and what one digs is one's *financial grave*. A single action is simultaneously investing and digging; a single condition is simultaneously having finished the digging and having lost one's money. Digging your own grave does not kill you, but digging your own financial grave does cause your death/bankruptcy.

Such blends can of course be elaborate, as in Seana Coulson's example from an editorial in the *UCSD Guardian*:

The U.S. is in a position to exhume itself from the shallow grave that we've dug for ourselves.

In this blend, the digger is identical to the body buried, which can exhume itself. This is impossible for the source Input, but possible for the target Input, where a nation can be in bad conditions but try to get out of them. In the blend, the ease of exhuming is related to the depth of the grave. This logic is available from both source and target Inputs: the shallower the grave, the easier the exhumation; the less bad the conditions, the easier it is to improve them. As in "you are digging your own grave," the actor is responsible but unaware, his actions were unwise, he is culpable for not recognizing that his actions were unwise, and the consequences of those actions are undesirable.

Pattern completion is at work in developing this blend. In recent U.S. history, there have been many disparate events, only some caused by actors, only some caused by American actors, and almost none caused by any single actor. Nonetheless, the blend asks us to integrate those many disparate target events, by blending them with a template, available to the blend from the source Input, of a single integrated action by a single actor, namely, digging as done by a digger. To do so, we must construct in the target a single entity, "the United States," that is causal for those many disparate events, which are in turn causal for current conditions in the United States. In the blend, the United States is a person, whom we want to convince to begin the process of self-exhumation.

Analogical counterfactuals

Consider an analogical counterfactual of the type studied by Fauconnier (1990, in press):

"In France, Watergate would not have harmed Nixon."

Uncontroversially, understanding this counterfactual includes building a generic space that fits both *American politics* and *French politics*. It includes a leader who is elected, who is a member of a political party, and who is constrained by laws. This skeletal generic space fits the space of *American politics* and *French politics* so well and intricately that it is natural for someone to project a great deal more skeletal information from *American politics* into the generic space on the assumption that it will of course apply to *French politics*.

The rhetorical motive for saying, "In France, Watergate wouldn't have done Nixon any harm" is exactly to stop someone from projecting certain kinds of information to the generic space on the assumption that it applies to *French politics*. The speaker lays down a limit to this projection by constructing a specific, counterfactual, and pragmatically anomalous blend.

Into this blend, the speaker has projected information associated with President Nixon and the Watergate break-in. Nixon and Watergate and so on are brought into the blend with only skeletal properties, such as being a president who breaks laws in order to place members of a political party at a disadvantage. It may be that such information in fact in no way belongs to *French politics*, that something like Watergate has in fact never happened in *French politics*. No matter, it can be imported to the blend from the "Nixon in America" input. Additionally, from the "France" input, we can project to the blend French cultural perspectives on such an event.

This counterfactual blended space operates according to its own logic. In this counterfactual blend, an illegal act directed with the knowledge of the elected leader against the opposing political party leader will not cause the public outrage associated with Watergate. For this central inference to take place, we must have both the nature of the event from the "Nixon in America" input and the general cultural attitudes from the "France"

input. The blend is not a side-show or curiosity or merely an entertaining excrescence of the projection. It is the engine of the central inferences.

The criterial properties of blending are apparent: cross-space mapping of the Input U.S. and France spaces; generic politics space; selective projection—Nixon and Watergate on the one hand, the frame of French politics on the other; emergent structure:

—composition provides a Watergate-like event in France;

—elaboration includes the explicit predication that the president is not harmed.

Finally, there is projection back to the Inputs: France has features that the U. S. does not have.

Clearly, in the case of such an analogical counterfactual, the construction of meaning cannot be mistaken as an attempt to impose structure from the one input onto the other. In fact, this particular analogical counterfactual is trying to do exactly the opposite. It is trying to make clear in just what areas information projected from one input cannot be imposed on the other. Moreover, its purpose is to illuminate not only the nature of the "France" Input, but also the nature of the "America" Input. The inferences are thus not one-way. They can go from the counterfactual space to both of the Inputs.

Nor are the analogical connections exclusively positive. It is disanalogy rather than analogy that is the central assertion of the statement. We recognize that a scenario can be shared by *American politics* and *French politics* but that in certain key respects these spaces have negative counterparts rather than positive counterparts. The utterance sets up a blend exactly for the purpose of illuminating these counterparts and their negative relation to each other. The projection in the case of "In France, Watergate would not have harmed Nixon," is thus not direct, not one-way, and not exclusively positive. This example lets us add to our model of conceptual projection the feature that even when, as in analogy, one input is in some way "understood" by projection from the other, the projection is in general not direct, not one-way, and not exclusively positive.

Of course, one may object to the assertion about France. One can respond, "You are wrong, look at all the harm the Greenpeace incident did to Mitterand." This can be interpreted as asking us to change the blend so that the illegal act is now general enough to include not only acts directed at an opposing political party but even acts directed against any opposing group (Greenpeace). It asserts that the space does indeed include cultural perspectives that, contrary to the previous assertion, do apply to both *American politics* and *French politics*. This, in turn, has the effect of expanding the generic space. This is a fundamental and general point that will arise repeatedly in our analyses: the array of spaces is built up dynamically and inventively in order to achieve a conceptual projection. The many-space model dictates no fixed sequence in this construction of meaning. It additionally accords notable place to energetic and imaginative effort and revision. It should

also be emphasized that while the English sentence *In France, Watergate ...* instructs us to perform a blend, it considerably underspecifies what blend to perform. There are countless other interpretations of this sentence corresponding to different blending choices (e.g., it could be about the love of the French for Nixon, or the consequences for Nixon of living in France rather than running for a second term, and so on). Rather remarkably, we are capable of constructing the "right" blend in context, in spite of the sparse grammatical clues.

We might ask, in what space does it hold that Watergate does not harm Nixon? Not in the "Nixon in America" Input, or the "France" Input, or in the Generic space. But if we shift to the blend, then the claim holds. It appears that a central part of conceptual projection is knowing how to construct a blend and how to shift to that blend in order to do real conceptual work, with the consequence that the vestiges of that real conceptual work are often projected to the one or both of the Inputs. But the structures of the blended space that would be impossible in the other spaces are left behind in such projection. That they are left behind does not mean that they are not indispensable to the central conceptual work.

Counterfactuals are not exotic curiosities of language. They are central to reasoning in everyday life (Kahneman 1995), and to scientific reasoning (Goodman 1947). Tetlock and Belkin (1996) show that argumentation in political science relies indispensably and extensively on counterfactual thought. Turner (1996a) shows that political scientists and others have not taken into account the complex blending that underlies the construction of counterfactuals, and the great range of conceptual structure and knowledge that it recruits without our noticing it (Turner & Fauconnier, in press). The biases smoothly integrated into the blend may serve the rhetorician, but not the social scientist.

Category extension and change

We frequently organize new material by extending a conventional category to it. Usually, these on-line category extensions are provisional, for local purposes, often purposes of expression and naming. Consider the attested case in which a handout for an academic talk has one column with elements listed 1 through 7, and another column with elements listed A through F. During the question period, people begin referring unselfconsciously to "Number E." The inputs to this blend are (1) the counting numbers and (2) the alphabet, ordered in its customary linear fashion. The generic space has only a well-ordered ordinal sequence. It defines the counterparts in the two Inputs. The blend has the well-ordered ordinal sequence, but also has, linked to it and thus to each other, two paired sets of counting numbers, one of which is the "real" counting numbers and the other of which is the alphabet. But the blend does not have, for example, arithmetic properties from the input space with counting numbers, or spelling from the space with the alphabet.

In other cases, the blend may lead to permanent category change. Consider the phrase "same-sex marriage". In Turner and Fauconnier (1995), we show in detail how expressions with this syntactic form can be systematically used to trigger blends. For *same-sex marriage*, the Inputs are the traditional scenario of marriage on the one hand, and an alternative domestic scenario involving two people of the same sex on the other. The cross-space mapping may link prototypical elements like partners, common dwellings,

commitment, love, etc. Selective projection then recruits additional structure from each Input, e.g. social recognition, wedding ceremonies, mode of taxation, and so on from the first Input of "traditional marriage," and same sex, no biologically common children, culturally defined roles of the partners, and so on from the second Input. Emergent properties will characterize this new social structure reflected by the blend.

At that stage of the construction, *same-sex marriage* will not be a subcategory of *marriage* for those who view *marriage* as having criterial attributes (e.g. heterosexual union for the sake of children) that *same-sex marriage* does not have. But now there can be pressure for these criterial attributes to change. The pressure comes from the activated generic space which made the blend possible. If that generic space (people living in a household, division of labor, mutual protection, financial planning done as a unit, or whatever) is understood to provide the essential criteria for the notion *marriage*, then *same-sex marriage* becomes a banal subcategory of the more general notion. Analogy and blending drive categorization. Clearly, different people using the same words in the same language may nevertheless entertain different categorization schemes. The same expression "same-sex marriage" may correspond to an analogical and conflictual blend for one person, and to a straightforward subcategory for another. Interestingly, the clashing conceptions of two such persons will still share a large amount of meaning.

Regatta

Let us consider another case in which it is clear that the motivation for constructing the blend is to tell us something about an important input. A modern catamaran *Great America II*, sailing from San Francisco to Boston in 1993, is being compared to a clipper, *Northern Light*, that made the same run back in 1853. A few days before the catamaran reached Boston, observers were able to say:

At this point, *Great America II* is 4.5 days ahead of *Northern Light*.

This expression frames the two boats as sailing on the same course during the same time period in 1993. It blends the event of 1853 and the event of 1993 into a single event. All the conditions for blending obtain. There is a cross-space mapping which links the two trajectories, the two boats, the two time periods, positions on the course, etc. Projection to the blend from the Inputs is partial: the 1853 date is dropped, as are the 1853 weather conditions, the purpose of the trip, and so on. But the blend has rich emergent structure: like the traveling monks, the boats are now in a position to be compared, so that one can be "ahead" of the other. This structure itself, two boats moving in the same direction on the same course and having departed from San Francisco on the same day, fits into an obvious and familiar cultural frame, that of a *race*. This yields additional emergent structure by completion. The race frame in the blended space may be invoked more explicitly, as in:

At this point, *Great America II* is barely maintaining a 4.5 day lead over *Northern Light*.

"Maintaining a lead" is an intentional part of a race. Although in reality the catamaran is sailing alone, and the clipper's run took place 140 years before, the situation is described in terms of the blended space, in which, so to speak, the two boats left San Francisco on the same day in 1993, and are engaged in a race to Boston. As in the monk example, no one is fooled by the blend: the clipper has not magically reappeared. The blend remains solidly linked to the Inputs. Inferences from the Blend can be projected back to the inputs: in particular, the speeds and positions of the two boats on their respective runs many years apart can be projected back to the inputs. Another noteworthy property of the *race* frame in the blend is its emotional content. Sailors in a race are driven by emotions linked to winning, leading, losing, gaining, and so forth. This emotional value can be projected to Input 2. The solitary run of *Great America II* is conceived, thanks to the blend, as a race against the nineteenth century clipper, and can be lived with corresponding emotions.

The attested report that prompted our interest in the "boat race" was actually a magazine article in *Latitude* 38, which contained the following:

As we went to press, Rich Wilson and Bill Biewenga were barely maintaining a 4.5 day lead over the ghost of the clipper Northern Light, ...

The blend, here, has become reified. An explicit referent, the ghost, is set up for the opponent of *Great America II* in the blended space. The mapping is more extensive, although still implicit. "*Ghost*" allows the projection from Input 1 that the clipper no longer (i.e. in 1993) exists. But the starting times are still fused, and it is understood that the "ghost" is retracing the exact run of the record-holding clipper.

Again, nobody is fooled into confusing the blend with reality. There is no inference that the sailors actually saw a ghost ship or even imagined one. The construction and operation of the blend is creative, but also conventional in the sense that readers know immediately and without conscious effort how to interpret it.

Because blending is neither deterministic nor compositional, there is more than one way to construct an acceptable blend, and this is confirmed by our boat race example. The preferred reading seems to be that 4.5 days is the difference between the time N it took *Great America II* to reach its current position (point A), and the time N+4.5 it took *Northern Light* back in 1853 to reach point A. Under that interpretation, the boats' positions in the initial spaces (1853, 1993), and in the blend, are their positions (point A for GA, and point B for NL) after N days, which is the time on the clock in the 1993 space at the time of writing. In this reading, the 4.5 days are a time in the 1853 space—the time it took NL to get from B to A. Another conceivable reading has this reversed, taking the time on the clock in the 1853 space and the 4.5 days in the current 1993 space. Under that interpretation, *Northern Light* got to point B' after N days, *Great America II* got to point A after N days, and it took *Great America II* 4.5 days to get from B' to A.

Other readings may be available. Suppose *Great America II* is following a different course from its illustrious predecessor's, so that positions on the two journeys cannot be directly compared. But suppose also that experts can estimate, given current positions, how long it "should" take *Great America II* to reach

Boston. Then, the example sentence could be interpreted as saying that, given its current position, *Great America II* should end up making the run to Boston in 76 days, 8 hours minus 4.5 days, i.e. in 71 days, 20 hours. This time, in the blended space of 1853 and the experts' hypothetical 1993 space, *Great America II* reaches Boston 4.5 days ahead of Northern Light.

All these readings involve blended spaces. The blended space is different in each case, and its structure accounts for the corresponding difference of truth values in the interpretations. This is a nice point: far from being fuzzy and fantastic, the blends allow a totally precise quantified evaluation of the truth conditions they impose on the actual world.

The desktop

Now take a superficially very different example, offered by Dan Gruen, which involves the performance of a specific activity. Human-computer interfaces are often structured by the concept of a desktop, on which objects rest and can be manipulated and used to perform actions. The appearance of the computer screen carries icons corresponding to objects on a desktop. They can be opened and closed, put away, and so on. When working with the icons, we think of them and act upon them in some ways as we would on actual desktop material, and in some ways as when dealing with general computer commands. Clearly, the entire activity is coherent and integrated, once learned. It is not hampered by its obvious literal falsities: there is no actual desk, no folders, no putting of objects into folders, no shuffling of objects from one folder to another, no putting of objects into the trash, and so on.

The created blend has considerable emergent structure. For instance, dragging icons with the mouse belongs to neither moving objects on a desktop nor giving standard symbolic commands, or *a fortiori* using the machine language. The user is not manipulating this computer interface by means of an elaborate conscious analogy, but as an integrated form with its own coherent structure and properties. From an "objective" point of view, this activity is totally novel—it shares no physical characteristics with moving real folders, and it is novel even for the traditional user of a computer who has issued commands exclusively from a keyboard rather than from a mouse. Yet the whole point of the desktop interface is that the integrated activity is immediately accessible and congenial. The reason, of course, is that a felicitous blend has been achieved which naturally inherits, in partial fashion, the right conceptual structure from both inputs, and then cultivates it into a fuller activity under pressure and constraints from reality and background knowledge.

The desktop example also nicely illustrates the non-arbitrary nature of blending: not just any discordant combination can be projected to the blend. Some discordant structure is irrelevant because it has no bad consequences—e.g., the trash can and the folders both sit on the desktop—but other discordant structure is objectionable—dragging the icon for a floppy disk to the trash as a command to eject the disk from the drive is notoriously disturbing to users. The inference from the domain of working at a desk that everything going into the trash is lost, and from the domain of computer use that everything deleted is irrecoverable, interfere with the intended inference that the trash can is a one-way chute between two worlds—the desktop interface and your actual desk.

Another point illustrated by the example is that input spaces are themselves often blends, often with an elaborate conceptual history. The domain of computer use has as input spaces, among possible others, the domain of computer operation and the domain of interpersonal command and performance. It is common to conceive of the deletion of files as an operation of complete destruction performed by the system at the command of the user. In fact, in the domain of actual computer operation, the files are not erased by that command, and can often be recovered. The user's sense of "deletion" is already a blend of computer operation and human activity. More generally, it is the fact that, by means of blending, keyboard manipulation is already conceived as simultaneously typing and high-level action and interaction that provides the appropriate partial structure to later blends like desktops with icons. The existence of a good blend can make possible the development of a better blend.

V. Advanced aspects of the network model

The previous sections have outlined the general characteristics of the cognitive operation of blending as reflected in superficially very diverse cases. Blending as a cognitive operation is elegant and uniform, but offers a great variety of different instantiations. A general program of research arises from inquiring into the general features of blending, the variety of purposes it serves, and the different ways in which it can be formally applied. In this section, we consider further general features of blending and constraints on the process. In subsequent sections, we consider more detailed taxonomies of blends according to structure, function, status with respect to reality, and internal logic. In all of these sections, we present research questions for the theory of blending and offer in some cases provisional or partial answers.

Spaces, domains, and knowledge.

A mental space is built up in part by recruiting structure from (possibly many) conceptual domains and from local context. We can build different and incompatible spaces by recruiting from the same conceptual domain. Consider a personification of death as an evil magician versus a benevolent magician: the evil magician makes objects disappear forever, while the benevolent magician transforms objects into other objects. The evil magician is a personification for a standard notion of death; the benevolent magician is a personification for a notion of death as involving reincarnation. In each case, one input space is built up by recruiting from the conceptual domain for *magician* and the other is built up by recruiting from the conceptual domain for *death*. But the two cases recruit different structure. The generic spaces have different event structure (deletion versus transformation). The blended spaces have different structure (evil versus benevolent magician). The feature of *evil* versus *benevolent* arises as an inference from blending—in the source conceptual domain of magic, there is nothing evil about making an object vanish and nothing benevolent about turning it into something else; but in the blend, the object is *us*, and our attitudes about our own vanishing or transformation provide the evaluations.

Consider also "Italian is the daughter of Latin" versus "Latin is the daughter of Italian [because students of Italian become interested in studying Latin]." Each has input spaces built up from conceptual domains of progeneration and languages, but quite different structure is recruited from the conceptual domains into the

input spaces. "Italian is the daughter of Latin: her ostentatious beauty is really a rebellion against her mother's austerity" recruits yet different structure from these conceptual domains to the input spaces. All three of these examples have the identical underlying conceptual domains, but quite different input spaces, generic spaces, and blends.

In the many-space model of conceptual projection, meaning is not constructed in any single space, but resides in the entire array and its connections. The "meaning" is not contained in the blended space. We know each space in the array and can work and modify all of them and their connections. During blending, conceptual work may be required at any site in the conceptual array. Spaces, domains, and frames can proliferate and be modified. Blending can be applied successively during that proliferation. Achieving useful counterpart structure and useful integration may require activating different input mental spaces, changing the recruitment of structure to them, establishing different generic connections between them, projecting different structure from the inputs to the blend, recruiting different frames to the blend, projecting different structure from the blend back to the inputs, multiplying the blends, and so on.

Integration of events


A fundamental motivating factor of blending is the integration of several events into a single unit. For example, although the boat race blend depends upon extensive connection of counterparts across different spaces, it also has integration of events: the sailing from one space is integrated with the sailing from the other space into a single event of racing, and this is the central point of the blend. In the desktop case, an action performed by the user of the computer is a single event that conceptually integrates the computer command and the manipulation of office items. It thus integrates both event components and conceptual counterparts. Even metaphoric mappings that ostensibly look most as if they depend entirely on the construction of metaphoric counterparts can have integration of events as a principal motivation and product. "He digested the book" of course has metaphoric counterparts, such as food and book, but it also projects an integration of events. In the source, digesting already constitutes an integration of a number of different events. But its counterpart in the target is, independent of the metaphor, a series of discrete events—taking up the book, reading it, parsing its individual sentences, finishing it, thinking about it, understanding it as a whole, and so on. The integrity in the source is projected to the blend so that this array of events in the target acquires a conceptual integration of its events into a unit. On one hand, the metaphor blends conceptual counterparts in the two spaces—eating and reading. On the other hand, the metaphor helps us to integrate some distinct event sequences in the space of reading. The blend exploits the integrity of events already present in the space of eating, and exports that integrity of events to the target space of reading. In the "digesting" metaphor, we export the integrity in the blend to induce an integrity of events in the target (picking up book, reading lines, finishing book, thinking about it, etc.). In the boat race, we export the integrity of events in the blend to induce an integrity of events in 1993 (preparing the boat, raising money, waving goodbye to well-wishers at the dock, trimming the sails, keeping the log, arriving at Boston, parting afterward, etc. etc. etc.) In both cases, there is a great range of events in one space (reading, 1993) that comes to acquire the integrity of an event structure in the blend (digesting, race). In some cases, like "digesting the book," the integration of events is already provided in one of the inputs and is recruited by the blend to provide integration for the other input. In other cases, like the boat race, the integration emerges in the blend.

In grammar, certain abstract scenarios are represented by corresponding grammatical constructions. A given construction goes with a given schematic scenario. To describe events using that construction is to prompt the hearer to integrate those events into that schematic scenario. "John kicked the ball over the fence" describes events of kicking and motion of the ball in a direction. It uses a construction that represents the schematic scenario in which an action causes an object to move in a direction. When we use the same construction to describe an act of praying and an event of boys coming home in "We prayed the boys home," [4] we are prompting hearers to integrate the events into the pattern of caused motion.

Recruiting and integrating internal connections from the inputs into the blend

Inputs will have internal connections that are motivated conceptually and experientially. For example, if the topic is a newspaper company, that company is linked to the newspaper (its product), the building (its location), its publicly-traded shares, and so on. As Nunberg (1978) has discussed, these connections motivate expressions like "The newspaper is on Main Street," "The newspaper went out of business," "The newspaper was sold for fifty million dollars," and so on.

Blends make use of these connections in several creative ways. Consider the following example of a cartoon representing a powerful newspaper company about to succeed in a hostile takeover of a weaker automobile company that will be eliminated by selling off its assets. The cartoon shows a giant printing press smashing a car. This is a metaphorical blend like those we have seen in section IV: input one has the stronger and weaker objects; input two has the contest between companies. The cross-space mapping is the basic metaphor that maps stronger objects destroying weaker objects to winning and losing. The strong heavy object is mapped onto the powerful newspaper company; the weaker object is mapped onto the weaker automobile company. But in the blend, we find the printing press as the strong heavy object and the car as the weak object. This is an efficient exploitation of internal connections: the printing press is a salient instrument of producing newspapers, and cars are the salient products of automobile companies. In the input, the printing press is not an instrument of destruction, but it has a force-dynamic function associated with crushing which can be associated with a car-smashing machine of the sort used in recycling automobiles. In the blend, the printing press is fused with both the company and the car-smashing machine. What is going on here? The blend must achieve three goals. First, given that the cartoon is a visual representation, the blend must be concrete and specific. Second, it must fit the frame of stronger and weaker object. Third, these objects in the blend must be properly connected to the companies in input two. The companies in input two, being abstract, cannot in themselves provide the corresponding concrete elements in the blend. The weaker and stronger objects in input one are concrete but not specific, and so cannot in themselves provide the corresponding specific elements in the blend. But we can exploit internal connections in the inputs to make the elements in the blend adequate. The printing press and the car are concrete, specific objects associated with the companies that can also be fit into the frame of the stronger object destroying the weaker object. They fit this frame in part because the printing press intrinsically has force-dynamic structure capable of destruction and in part because we are familiar with car-smashing machines. In the blend, two elements are simultaneously (1) two concrete, specific objects; (2) a stronger object destroying a weaker object; and (3) two companies.



Clearly, such a blend is creative. Not just any connections will do. There has to be a search for elements that simultaneously satisfy a number of constraints. Below, we will discuss some candidate constraints for recruiting internal connections to blends.

Opportunism and path-dependency

Although the laws of biology motivate all biological change, it is not possible to predict the evolution of a species, since its evolution will depend at each step on local accidents. The genetic structure that evolution has to work with at any moment depends upon the history of those accidents. The path of accidents shows opportunistic exploitation of existing structures of the organism or features of the environment. We can speak because an existing mammalian supralaryngeal airway, previously adapted for breathing and eating, could be opportunistically adapted for speech, at the cost, as Darwin observed, of making us liable to choke to death on our food.

Similarly, blending shows us that reason looks for accidents to exploit opportunistically. It is accidental that fusing the monk's paths and days but not the monks results in a blend that is easily completed by our standard frame of *two people approaching each other along a path*, but this serendipitous accident, once found, provides the solution to the problem. The printing-press blend is effective only because we know about printing presses and about car-recycling machines that happen to share a force-dynamic structure with printing presses. Had historical accident given us a world without these car-recycling machines (e.g., the world in 1950) and with a different prototypical method of printing (e.g., spraying ink), an entirely different blend would be required in order to achieve specificity, concreteness, conformity to the frame of stronger and weaker object, and proper connection to the companies in input two. Constructing that blend would require opportunism in the seeking of accidents to exploit.

Opportunism is sometimes displayed as a mark of wit: Consider "Banging a Tin Cup With a Silver Spoon." This headline announced a news story about Orange County, whose financial managers lost much of the county's assets betting heavily on interest trends. Although the county remained extremely rich, it declared bankruptcy and asked creditors for debt forgiveness. The reporter described the county as a "wealthy deadbeat." In one blend, Orange County is personified as a beggar with a tin cup. In another, it is personified as a wealthy individual with a silver spoon. Both of these blends are conventional. In a hyperblend of these two blends, the county is personified as a wealthy beggar. It is accidental that a person can hold both a tin cup and a silver spoon and bang the cup with the spoon in the manner of a beggar drawing attention to his begging. The headline asks for applause for its ingenuity in finding these accidental connections. This turns out to be a general property of blends: they are judged to be better according as they exploit more accidental connections.

Entrenchment

Like other forms of thought and action, blends can be either entrenched or novel. "Digging your own grave" is a complex blend entrenched conceptually and linguistically. The Buddhist monk blend is novel and is used for only that one riddle. We often recruit entrenched projections to help us do on-line conceptual projection. On-line projections and entrenched projections are not different in kind; entrenched projections are on-line projections that have become entrenched. Our seemingly fixed projections are highly entrenched projections of an imaginative sort. Because the mechanisms of projection are shared in the two cases, entrenched structures are subject to transformation under work by on-line projection.

Fusion

Fusion of Counterparts

Blending can fuse counterparts in input spaces. In the monk example, the days are fused and the positions are fused; in the debate with Kant, fusion operates over issues, languages used, and modes of expression for debate.

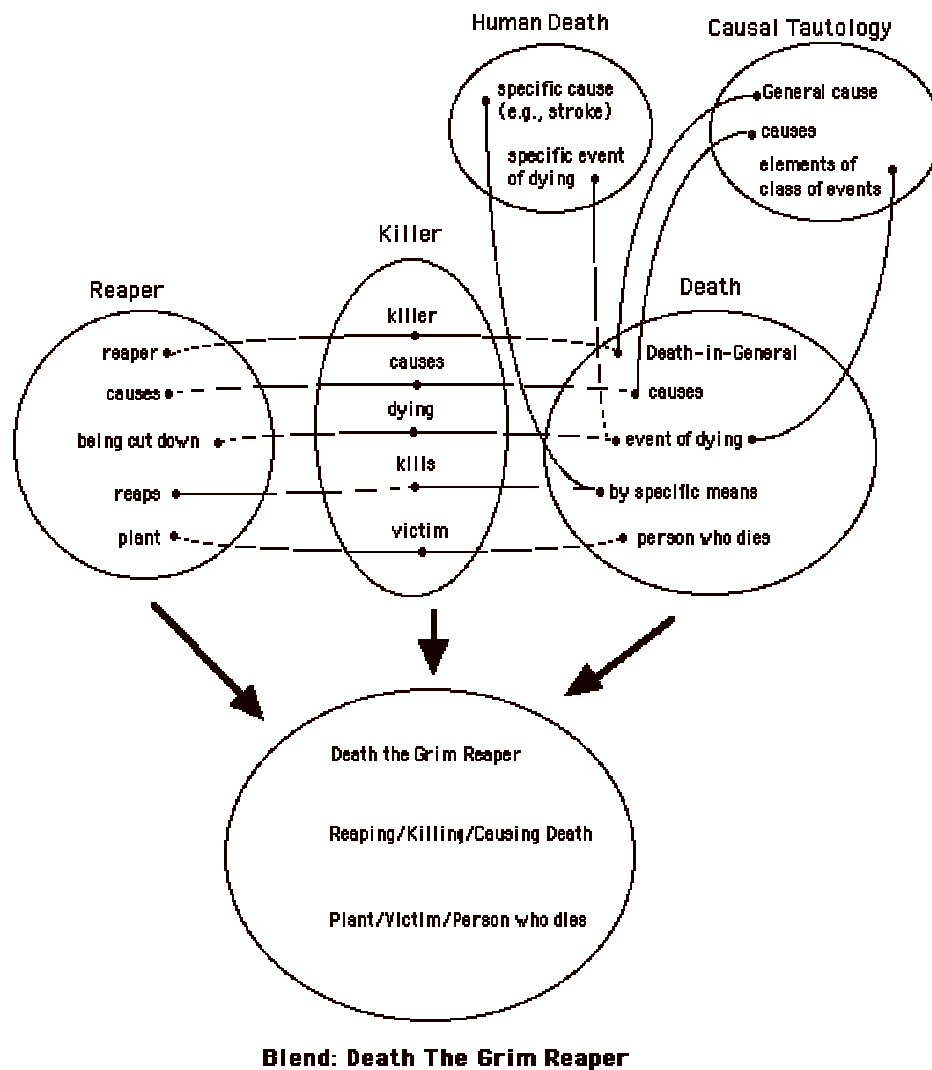
But the fusion is not always simple. In the debate blend, the time of the debate is a fusion—there is only one time in the blend, not two times. But it is neither the time of the inputs nor some combination of them. It is a special transcendent time—it would be odd to say, "Two years ago, Kant disagreed with me, when I thought reason was a self-developing capacity."

Non-fusion of counterparts

Blending need not fuse counterparts in input spaces. In the monk example, the two monks are not fused. In regatta, the two boats are not fused. In the debate with Kant, Kant and the modern philosopher are not fused.

Combination of non-counterparts

Blends can combine non-counterpart elements that come from different inputs. Consider The Grim Reaper, which is a blend with several input spaces, including a space of harvest and a space of particular human death. A reaper in input 1 is the counterpart of Death in input 2, not of the skeleton, but since Death as a cause is metonymically associated with *skeleton* as an effect, the blend can combine the reaper (from one input) with the skeleton (from the other), even though they are not counterparts. Similarly, elements in a single input space that are metonymically related can be combined in the blend. Priests, monks, mourners, and members of lay brotherhoods that are associated with dying, funerals, burial, and afterlife are metonymically associated with Death. They are not counterparts of Death, but in the blend, an attire we associate with them—robe and cowl—can be the attire of The Grim Reaper. The cowl, pulled over the head of The Grim Reaper, at once evokes both connotations of death and the impression of Death as mysterious, unknown, and set apart from human society.



The possibility of combining non-counterparts on the basis of metonymic connections—like the connection between Death and a skeleton—gives blending a great power: the blend can combine elements that contribute to the desired effect *even though those elements are not counterparts*. The combined elements "go together" in evoking the same effect even if they do not "go together" according to the counterpart connections between the input spaces.

Combination in the blend of non-counterparts is routine. In "He was red-hot with anger; I could see smoke coming out his ears," *heat* in one input has the metaphoric counterpart *anger* in the other input, but *anger* has a metonymic connection to physiological reactions, including redness of skin and increased body heat. *Heat* in the blend combines *heat* from the source input, *anger* from the target input, and *body heat* from the target input, even though the two "heats" in the inputs are not counterparts in the metaphor. (See the study by Lakoff and Kövecses, described in Lakoff, 1987.) The Birth Stork blend, which is based on the counterparts provided by the conventional metaphor BIRTH IS ARRIVAL, ingeniously provides a stork-with-diaper-sling that

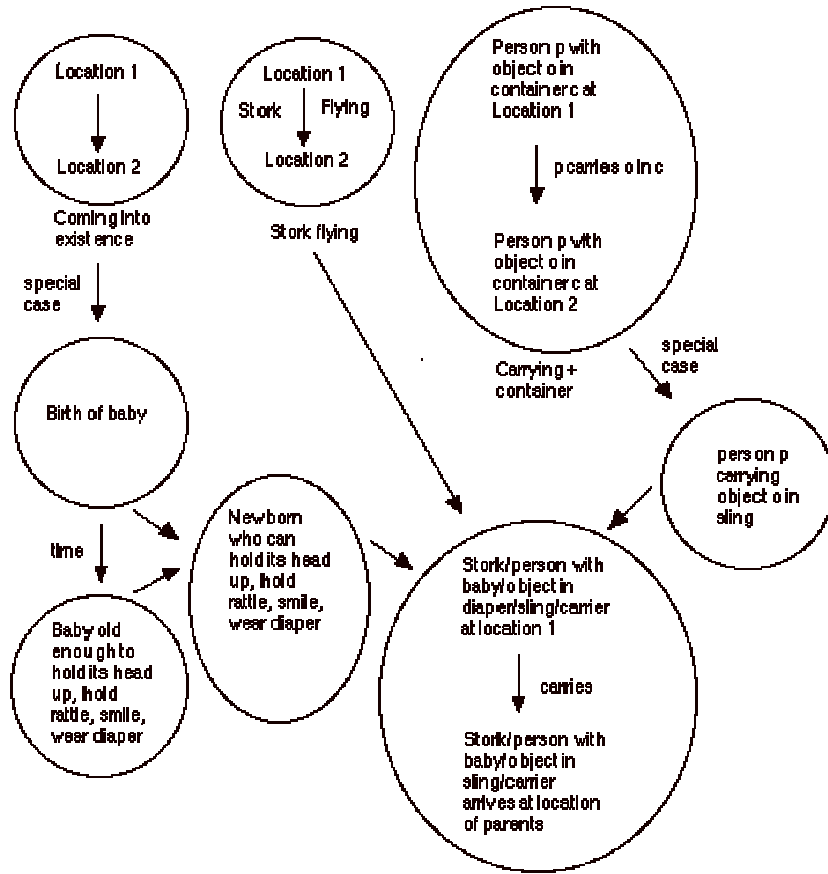
has as its counterpart in one input the vehicle of ARRIVAL and in the other input general causal processes of birth; the diaper belongs to neither of these counterparts, but because the baby, which is the product of birth, is metonymically associated with diapers, the diaper can be combined with the general process of birth and used concretely as part of the blend-vehicle.

The Birth Stork network makes use of some pre-existing blends. When an element in one state is later in a different state, we can compress this into a space in which the element undergoes a “change of state.” When an element in one location is later in a different location, we can compress this into a space in which the element undergoes a “change of location.” In general, when two spaces are related by both counterfactuality and temporal distance, we have the chance to compress those spaces and their vital relations into a single “change” blend. These two networks, “change of state” and “change of location” have, as metaphor theorists have noted, served as inputs to a further blend in which the change of state is blended with the change of location, as in “the water is coming to a boil.”

There is a third network related to these two: in one space, there is an element, but in a temporally prior space, there is no element. These two spaces also have outer-space vital relations of counterfactuality and time, and, following the general pattern, they are compressed into a single blend in which the element is always there, but undergoes a change of state from non-existence to existence.

This third blend of nothing-then-something is blended with the standard change of state/location blend into a very standard blend of “coming into existence.” In this way, we understand the passage from nothing to something as motion from one location to another. This “coming into existence” blend is naturally used to frame birth as the coming into existence of the baby: “Has the baby arrived yet?” “It’s on its way.” “It should be here any day now.”

The following diagram takes as one of its inputs this already complicated blend of “coming into existence.”



Biases

Composition, completion, and elaboration all recruit selectively from our most favored patterns of knowing and thinking. This makes blending a powerful cognitive instrument, but it also makes it highly subject to bias. Composition, completion, and elaboration operate for the most part automatically and below the horizon of conscious observation. This makes the detection of biases difficult. Seepage into the blend can come from defaults, prototypes, category information, conventional scenarios, and any other routine knowledge.

VI. Optimality principles

Our discussion so far has shown a common structure for all blends, yet not all phenomena that manifest that structure are equally good blends. Some blends are better than others. There are optimality principles that a blend can meet more or less well. It turns out that these principles compete, and that there is a variety of more specific blend structures which are motivated by these principles. Here we discuss the principles we have been able to substantiate and some of these more specific blend structures.

Integration:

The blend must constitute a tightly integrated scene that can be manipulated as a unit. More generally, every space in the blend structure should have integration.

Topology:

For any input space and any element in that space projected into the blend, it is optimal for the relations of the element in the blend to match the relations of its counterpart.

Web:

Manipulating the blend as a unit must maintain the web of appropriate connections to the input spaces easily and without additional surveillance or computation.

Unpacking:

The blend alone must enable the understander to unpack the blend to reconstruct the inputs, the cross-space mapping, the generic space, and the network of connections between all these spaces

Good reason:

All things being equal, if an element appears in the blend, there will be pressure to find significance for this element. Significance will include relevant links to other spaces and relevant functions in running the blend.

Satisfaction of the optimality principles in some basic kinds of conceptual integration network.

Frame networks. To see a standard strategy of satisfying these optimality principles, consider again three examples: the Buddhist monk, the Debate, and Regatta. Of course, they all have cross-space mapping, selective projection to the blend, and a generic space that applies to both inputs. In addition, in each of these cases, all of the spaces share a rich frame and much of its content: in the Buddhist monk, all the spaces have *man walking along a mountain path*; in the Debate, all the spaces have *philosopher musing on a philosophical problem*; in Regatta, all the spaces have *boat sailing along an ocean course*.

A *frame network* is a conceptual integration network in which all spaces, inputs, generic, and blend, share topology given by an organizing frame. An organizing frame for a mental space is a frame that specifies the nature of the relevant activity, events, and participants. An abstract frame like *competition* is not an organizing frame, because it does not specify a cognitively representable type of activity and event structure.

Regatta, Debate with Kant, and the Buddhist monk are all frame networks. Typically, in a frame network, the common frame *F* inheres in the more elaborate frame *FB* in the blend. In the boat race example, the shared organizing frame *boat sailing along an ocean course* inheres in the more elaborate frame in the blend of *sailboats racing along an ocean course*. In the debate with Kant, the shared organizing frame *philosopher musing on a problem* inheres in the more elaborate frame in the blend of *philosophers debating about a problem*. In the Buddhist monk, the shared organizing frame *man walking along a mountain path* inheres in the more elaborate frame in the blend of *two men meeting on a mountain path*.

An organizing frame provides a topology for the space it organizes—that is, it provides a set of organizing relations among the elements in the space. When two spaces share the same organizing frame, they share the corresponding topology and so can easily be put into correspondence. Establishing a cross-space mapping between inputs is straightforward when they share the same organizing frame.


While spaces in a frame network share topology at the level of an organizing frame (we call this TF topology), they may differ at a more specific level (TS topology). For example, in the boat race network, there are two elements that fit the role *boat* in the organizing frame and so have identical TF topology. More specific relations, however, define finer topologies that often differ. For example, in the boat race network, one of the elements fits the more specific frame *nineteenth-century clipper on a freight run* and the other fits the more specific frame *late-twentieth-century exotic catamaran on a speed run*. The two more specific frames are different, and so the topologies are different at the TS level. More precisely, we reserve the term TS for finer topology that specifies values of roles that are in the organizing frame. These values may themselves be roles of a more finely specified frame. In our example, *boat* is a role of the organizing frame, *clipper* gives that role a more specific value and is itself a role of a more specific frame *clipper on a freight run*. Features of these more specific values—like monohull versus catamaran—can be projected to the blend.

There will also be incidental topology, TI, in both input spaces. We use the term TI for finer topology that

does not have to be included or specified, given the organizing frame. In our example, it may be fully part of the actual ocean voyages that dolphins escort the boats and that they pass by a certain uncharted island, but these are not assigned a role in the organizing frame. In general, features of incidental topology can also be projected to the blend.

The selection of an organizing frame for a space is not a once-and-for-all decision. The organizing frame can be modified and elaborated as the integration network is constructed. Topology at the TF, TS, or TI level may come to be promoted or demoted as needed. For example, *obstacles* may be a role in the frame *boat making an ocean voyage*, and if the clipper has difficulty traveling near the uncharted island because of technical problems of navigation that had not been solved in 1853 while the catamaran has difficulty traveling near dolphins because it is forbidden under international law from sailing through a school of dolphins, then *uncharted island* for the clipper and *dolphins* for the catamaran get promoted to the TS level as fitting the role *obstacle*, while *uncharted island* for the catamaran and *dolphins* for the clipper remain at the TI level.

Organizing frame is shared by all spaces:



Frame Network

Figure 8

Shared topology network. In a conceptual integration network over two inputs, the topology of the generic space is always shared by all four spaces—the blend, the two inputs, and the generic space. We will call a structure in which all spaces share the topology of a generic space a *shared topology network*. Four-space blends are *shared topology networks*, but multiple blends need not be, as we shall see below.

A frame network is a shared topology network whose shared topology is moreover an organizing frame. Other shared topology networks do not share organizing frame but do share topology. For example, simple

metaphors such as the portrayal of two business competitors as boxing opponents do not have a shared organizing frame: the source input in the example has *boxing* as its organizing frame, while the target input has *business competition* as its organizing frame. But the source and target inputs do share a higher-level structure of *competition* which gives them a shared topology and makes the cross-space mapping and the generic space possible.

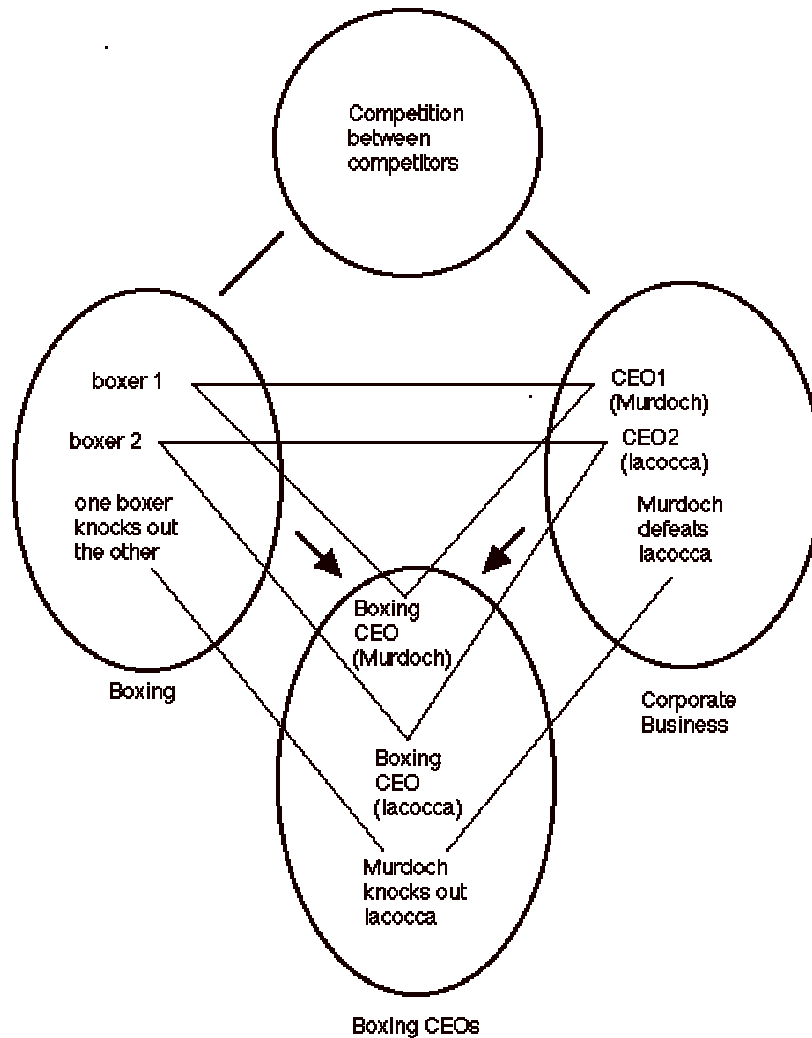
The case of complex numbers is another example, where one input has the organizing frame of two-dimensional geometry and the other has the very different organizing frame of real/imaginary numbers. The development of the cross-space mapping and the recognition of the topology shared by the inputs required a long and arduous period of conceptual work by mathematicians, and it was only at the end of the historical process that the generic space defined by this cross-space mapping came to be recognized and named: a commutative ring.

We can now define a *frame network* briefly and more systematically as a *shared topology network* whose generic space, cross-space mapping, and shared topology are all given by virtue of a shared organizing frame for all spaces.

One-sided networks. A shared topology network is one-sided if the inputs have different organizing frames and one of them is projected to organize the blend. Its defining property is that the organizing frame of the blend is an extension of the organizing frame of one of the inputs but not the other: $TFB > TF1$.

The case of the two boxing business competitors is a one-sided network, whose generic space has an abstract relation of adversarial competition between two agents. The blend inherits the frame of Input 1, *boxing*. The cross-space mapping is metaphoric, with Input 1 as the source and Input 2 as the target.

In a simple metaphoric blend like this, projection from inputs to blend is highly asymmetric: one of the inputs but not the other supplies the organizing frame and therefore frame-topology. This is why it seems appropriate to call that input the *source input*. The projection of the source frame to the blend carries with it linguistic constructions (e.g., vocabulary) used to evoke the source frame. Of course, there are projections from the target input to the blend that also provide linguistic constructions for the blend, but they refer to elements below the TF level, at the TS or TI level. For example, if the two business competitors are named Murdoch and Iacocca, we may say that "Murdoch knocked Iacocca out": "knocked out" belongs to the TF level of the source while "Murdoch" and "Iacocca" belong to the TS level of the target.



Any particular simple metaphoric one-sided network may have inhering within it a higher-order conventional metaphoric mapping, called by Lakoff and Johnson (1980) a *basic metaphor*. Such a basic metaphor is highly productive and inheres in indefinitely many particular constructions of meaning but is itself abstract. For example, the blend structure for the boxing business competitors is an active, on-line, specific conceptual structure that has inhering within it the abstract, basic metaphor of competition as physical combat. A basic metaphor itself never constitutes an active, complete, on-line construction of meaning. It always requires additional conceptual specification and projection to supply a particular construction of meaning.

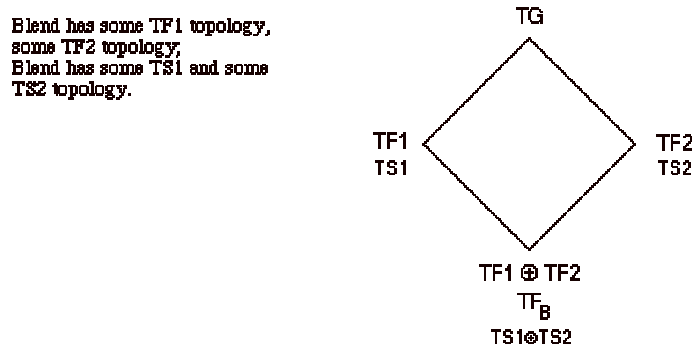
Two-sided networks. A shared topology network is two-sided if the inputs are organized by different frames but some topology is projected from both frames to the blend. Gruen's example of the computer desktop interface is a two-sided network. The two principal inputs have different organizing frames, the frame F1 of office work with folders, files, trashcans on one hand, and the frame F2 of traditional computer commands on the other. (There is also the lesser input of *choosing from a list*.) In the blend, some of the elements have F1 topology from one input while others have F2 topology from the other input.

The metaphor "digging your own grave" is also a two-sided network with frame structure projected from both inputs. Death and graves come from the source input of the "dying" scenario, but causality and intentionality are projected from the target input of discretionary action and mistakes that lead to failure, in the following way. In the target input, making mistakes is unintentional and brings one closer to failure. The blend receives this causal and intentional structure by selective projection from the target input: in the blend, digging is unintentional and brings one closer to death. But in the source input, both the causal order and the intentionality have the reverse structure: in this source, it is someone's dying that causes the grave to be dug and the digging is moreover intentional. The temporal order of events in the blend (digging before dying, making mistakes before failing) is also taken from the target input, not the source input.

Complex numbers are another case of a two-sided network. The inputs are respectively two dimensional space and real/imaginary numbers. Frame structure is projected from each of the inputs, e.g., angles, rotations, and coordinates from two-dimensional space, and multiplication, addition, and square roots, from the space of numbers.

We also see a two-sided network in *same-sex marriage*: input 1 has marriage but not same-sex partners; input 2 has same-sex partners but not marriage. The blend takes marriage from the TF level of input 1 and same-sex from the TF level of input 2.

In all these cases, as in most networks, the blended space develops emergent structure of its own, and ends up with a richer specific frame FB. For example, in the case of complex numbers, multiplication in the blend includes addition of angles. This operation is unavailable in either of the inputs. The input of two-dimensional space doesn't have multiplication; the input of numbers doesn't have angles.



Two-sided Shared Topology Network

Figure 9

Asymmetric two-sided networks. In an asymmetric two-sided shared topology network, although the blend receives projection of organizing frame-level topology from *both* inputs, nonetheless the organizing frame of the blend is an extension of the organizing frame of only one input.

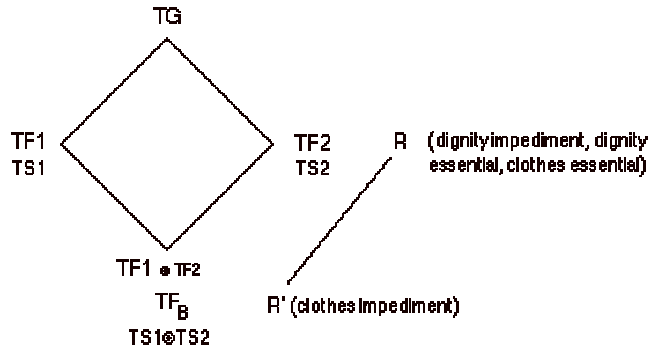
Consider as an example the case in which one person, observing that the Vatican seems to be flat-footed in the metaphorical boxing match over abortion, says, "I suppose it's hard to bob and weave when you have a mitre on your head." The Pope's competition with an adversary is portrayed as a boxing match, where the Pope is impeded as a boxer by the mitre he is obliged as Pope to wear on ritual occasions, and we interpret this as meaning (with respect to the input space with the Pope) that his obligation as Pope to remain dignified impedes him in his competition. In the input space with the Pope, there is a relationship at the level of the organizing frame between the Pope and dignified behavior and also between the Pope and his mitre. The cross-space mapping between inputs does not give counterparts in Input 1 for the *required dignity* or *required headgear* elements in Input 2. The Pope's obligation and his headgear in Input 2 both project to the headgear of the boxing Pope in the blend.

In the organizing frame of the input of boxing, the boxers have no headgear that is an impediment. In the blend, the organizing frame is slightly different: it contains the role *heavy headgear that makes fighting difficult*. This organizing frame is an extension of the frame of boxing, not of the frame of Pope and Roman Catholicism. Specifically, the frame of the blend has all the roles of *boxing*. But, the headgear—namely, the mitre—is projected from input 2. In that input 2 frame, there is a crucial relation R: the dignity of the Pope makes it harder for him to compete because he must always be honest and decorous. In input 2, the role *mitre* is directly linked (as a symbol) to the role *dignity and obligation of Pope*. The crucial relation R in input 2 is projected to R' in the blend: the mitre/dignity makes it harder for the Pope to box. *Mitre* and *dignity* in input 2 are both projected to the same element in the blend, and, crucially, they have no counterpart in input 1. The blend gets an organizing frame from input 1 but also the frame-level relation R from input 2, and this is what makes it two-sided.

In the blend, we find all the elements of the frame of boxing plus the heavy and unwieldy mitre on the boxer's head. It turns out that having a heavy object on the head is an impediment to fighting, and so we have a very natural and automatic pattern completion of the blend, leading to a new frame *boxing as impeded by heavy headgear*. This frame is an extension of the organizing frame of input 1, not of input 2, and this is what makes it asymmetric.

Recall that in *digging your own grave*, the cross-space mapping connected incompatible counterpart relations, such as direction of causality, and that to project causal direction to the blend, it was necessary therefore to choose one rather than the other of these counterpart relations. In the Pope example, because relation R in input 2 has no counterpart relation in input 1 (and *a fortiori* no incompatible counterpart relation), it can be projected to the blend (appropriately extended by completion), and no choice needs to be made between incompatible counterpart relations.

F_B is an extension of F_1 , but structure R from F_2 is also projected into that extension.



Asymmetric two-sided Network

Figure 10

Unfilled Shared Topology Networks. All conceptual projection appears to be particularly sensitive to certain kinds of abstract structure: causal relations, image-schematic relations, modalities, basic ontological categories, and event-shape. There are cases in shared topology networks where a relation is unspecified in the generic space but specified in incompatible ways in the inputs. That relation is *unfilled* in the generic space. We will say that the generic space is "partially unfilled" and that the shared topology network is *unfilled*.

In the frame network of the Buddhist monk, all four spaces share the frame *man walking along a (directed) mountain path*, but Input 1 has a direction for that motion (up), and Input 2 has a different direction for that motion (down). The specific direction of the motion is part of the event-shape of the motion, and moreover it is projected from each input into the blend. This does not create a clash in the blend, because the counterpart monks in the inputs are not fused when projected to the blend, so we have in the blend one monk ascending and the other descending. The two specific directions do not correspond in the inputs and are not connected in the cross-space mapping. The generic space is therefore partially unfilled, because it does not specify the direction of motion.

Two-sided unfilled shared topology networks. A two-sided shared topology network is often unfilled. *Digging your own grave* is a clear example. In the generic space, there is action by some agent and a bad event for some agent, and a causal relation between them. But the direction of causality is unspecified. In input 1, the bad event is death, the action is digging the grave, and the causality is from the bad event to the action: when someone dies, someone digs a grave. In input 2, the bad event is failure, the action is failing, and the causality is from the action to the bad event failure. The causal relation is in the generic space, but the specific direction of the causal relation is *unfilled*.

Single-framing networks. A basic kind of conceptual integration network we have not explored in this article is a single framing network. Briefly, a single framing network has an abstract frame as one input and as the other input a specific situation that has no organizing frame at all for the purpose of the integration, and so no

potential for competition with the organizing frame of the first input. For example, "Sally is the daughter of Paul" has the kinship frame *daughter-ego* as one input and as the other input a specific situation containing nothing but Sally and Paul. In the blend, Sally is framed as *daughter* and Paul is framed as *ego*. There is crucial emergent structure in this blend: the blend has a role *daughter of Paul* that is unavailable from either input. Moreover, the *ego* role in the kinship input is specified in the blend to be a father rather than a mother.

Metonymy

Before we complete the taxonomy of blend structures by topology, we must present a further optimality constraint concerning the projection of metonymies from inputs.

Recruiting special connections in one of the inputs can be used to bring in additional structure that assists in satisfying the optimality principles. Where an element in the blend has a topology that does not match the topology of a counterpart in one input, special connections internal to that input can be recruited to increase topological connections and help satisfy other optimality principles.

For example, in the Birth Stork blend, the diaper has a topology of being used as a sling (and more generally as part of the vehicle) which does not match the topology of the diaper in the Newborn input. But bringing the diaper into the blend helps satisfy Web, since it establishes more connections to the Birth space; and Unpacking, since it includes an element of the Birth space along with elements from the space of Transport; and Good Reason, since an alternative way of carrying the baby—in a paper sack, or net, for example—would not have the Good Reason of close association with the baby in the Birth space.

The analysis is similar for the cowl worn by the Grim Reaper. In the blend, the cowl has a topology—attire of the agent—that is not matched by the cowl in the input space. But exploiting the special and distant connection in the Death space of the relationship between dying and the priest and the attire of the priest, and thereby bringing the cowl into the blend, helps satisfy Web, Unpacking, and Good Reason.

The skeleton as the form of the Grim Reaper is slightly more complicated, because in this case the skeleton in the Death input has some useful topology—on its own, it is frightening or at least impressive as a salient result of the death of a human being. Exploiting the special connection in the Death input of the relationship between dying and the final result of a skeleton, in order to bring the skeleton into the blend, helps satisfy Topology as well.

We have seen a continuum of blend structures—from frame networks to one-sided, two-sided, and unfilled shared topology networks. Along this continuum, the topological connections in the basic cross-space

mapping between input spaces typically grow weaker, but other connections are employed for the purpose of maximizing Topology, Integration, and Web. By contrast, as this continuum is ascended, it grows easier to satisfy Unpacking since the blend increasingly incorporates special connections from one input without counterparts in the basic cross-space mapping. Generally, along this continuum, as the basic cross-space topology is weakened, Unpacking is strengthened.

The data for metonymy in blend structures provide evidence for an optimality constraint which we call the *metonymy projection constraint*:

Metonymy projection constraint:

When an element is projected from an input to the blend and a second element from that input is projected because of its metonymic link to the first, shorten the metonymic distance between them in the blend.

We saw above that blending can combine non-counterpart elements from a single input, such as Death, the cowl of the priest, and the skeleton of the person who has died. The metonymic distance is large between abstract death as the general cause of all deaths and the cowl worn by a certain kind of participant in a ritual associated with particular deaths. But in the blend, the metonymic connection is direct: the cowl is the attire of Death. Similarly, the skeleton after decomposition of the body is a distant product of death. But in the blend the skeleton is actually a body part of Death. The fact that metonymy is preserved in such cases can be viewed as a consequence of Topology. The metonymy projection constraint additionally specifies that metonymies get tighter under projection.

Satisfying the metonymic projection constraint is not a matter of blindly projecting metonymies. The internal integration of the blend provides opportunities for some acceptable metonymies but not for others. Since Death is an active person in the blend, and active persons are known to have skeletons (although they are not normally visible), the part-whole metonymy skeleton-body becomes available as the counterpart of the distant metonymy in the input.

Tightening metonymies under projection typically optimizes integration in the blend, since it helps build a tighter and more easily manipulated unit.

Now we return to the taxonomy of blend structures by topology.

One-sided shared topology network with metonymy projection. Suppose the example of the boxing business competitors is elaborated slightly—the competitors are now a newspaper magnate and an automobile magnate, and they are identifiable in part because one has a rolled-up newspaper in his back pocket and the

other has a car key on a key ring hanging out of his back pocket, each with an appropriate label. The organizing frame of the blend is still projected from the Boxing input, so the network is one-sided. But there is a frame relation in Input 2 that, in accord with the metonymy projection constraint, is projected to the TI level of the blend. The frame-relation in Input 2 is that the newspaper is the commercial product of the magnate's activities. The newspaper in the blend is connected to a newspaper in Input 2. The newspaper in Input 2 has no counterpart in Input 1 and its relevant topology in Input 2—product of the magnate's activities—is not the topology in the blend—copy of newspaper read by the boxer-magnate. The blend has an element—newspaper—projected from an input but the topology of that element in the blend is inherited from neither input. The metonymy between the magnate and the newspaper as commercial product in the blend is tightened under projection, so that it becomes part of the magnate's appearance. The analysis is similar for the car key.

Two-sided shared topology network, symmetric with metonymy projection. Recall the visual cartoon in which the printing press smashes the car. We pointed out that the printing press and car have topology in the blend (the press crushes and the car is crushed) that their counterparts in Input 2 do not have (the press is an instrument of making newspapers and the car is a salient product of the automobile company). Additionally, the printing press and car in Input 2 have no counterparts in Input 1. Interestingly, the elements that did not project their input-topology (printing press and car) end up being the only objects in the blend. This contrasts with the cartoon where the newspaper-in-the-back-pocket is only an optional element in the frame organizing the blend. The cartoon of the printing press smashing the car is remarkable because it is a case where Integration and Topology are maximized by recruiting special internal connections in Input 2. Because the topologies of strong and weak object on the one hand and competing companies on the other will match only at a very abstract level, we find that in addition to the companies, objects closely connected to them are projected to the blend in a way that closely matches and elaborates the Input 1 topology of strong and weak objects.

This blend structure is two-sided because the topology of strong and weak object comes from Input 1 but the topology of intentionality (the printing press intends to crush the car and the car hates it) comes from Input 2, where it is attached not to the printing press and the automobile but rather to the respective companies. The projection to the printing press and the car in the blend is symmetric: their topology in the blend matches frame topology in both inputs.

This example emphasizes that conceptual projection is a dynamic process that cannot be adequately represented by a static drawing. Once the conceptual projection is achieved, it may look as if the printing press has always corresponded to the stronger object and the car to the weaker. But in the cross-space mapping, the printing press and the car play no role; they have no counterparts in Input 1. Rather, the cross-space counterparts are stronger object and newspaper company, weaker object and automobile company. Under metonymy projection from Input 2, the printing press *in the blend* becomes the counterpart of the stronger object in Input 1, and the car *in the blend* becomes the counterpart of the weaker object in Input 1.

This example also shows that identity is metonymy of zero distance. The metonymic relation in Input 2 between company and commercial product is transformed into identity in the blend, where the printing press

is identically both a printing press and the newspaper company to which it is metonymically related as an instrument (in one of the inputs).

Two-sided shared topology network, symmetric with metonymy projection and additional frame recruitment. Suppose the cartoon now contains the newspaper magnate operating the printing press to smash the car, which is being driven by the car magnate. Here the blend structure becomes elaborate through the recruitment to the blend of an additional adversaries-with-instruments frame in which adversaries fight with opposing instruments, and in which the winning adversary has the superior instrument. Now the printing press and car in Input 2 have counterparts in the adversaries-with-instruments frame: in input 2, the printing press is a symbol of a capacity for productivity that is an instrument of corporate competition, and the car is a product that is an instrument of corporate competition; these instruments in Input 2 are the counterparts of the instruments in the adversaries-with-instruments frame. Now, the topology of opposing instruments in the blend matches the topology of opposing instruments in the adversaries-with-instruments frame. This frame has the useful property of aligning superiority of instrument with superiority of adversary. In this case, we see that exploiting special internal connections in Input 2 makes it possible to recruit a frame that makes Topology much stronger in the blend structure.

Optimality principles and one-sided networks

In the one-sided network exemplified by the business competitors portrayed metaphorically as boxers, Integration in the blend is automatically satisfied because the blend inherits an organizing frame from the source, *boxing*. Topology is satisfied between blend and source for the same reason. But Topology is also satisfied between blend and target because the conventional metaphor of competition as physical combat has aligned the relevant topologies of the source and target input spaces. Thus, when an element in the blend inherits topology from an element in either input that is involved in the cross-space metaphoric mapping, the topology it inherits is automatically, by virtue of the metaphor, compatible with the topology of that input element's counterpart in the other input. Web is similarly satisfied by this shared topology. Unpacking is provided just as it was for a frame network—although the blend is integrated at the TF level, it is disintegrated at the TS level. Suppose, for example, that the competitors are represented in a cartoon as boxing in business suits. This lack of integration between business suits and boxing prompts us to unpack to two different spaces, one of boxing and one of business. In the same way, if we know that "Murdoch" and "Iacocca" refer to businessmen and not boxers, then their use in the sentence "Murdoch knocked out Iacocca" directs us to the TS level of the input of businessmen, and this helps satisfy Unpacking.

Optimality principles and two-sided networks

In a two-sided network, Topology, Integration, and Web are not satisfied in such an automatic and routine fashion: it is necessary to use a frame that has been developed specifically for the blend and that has central emergent structure. (This may be why two-sided networks—such as the desktop, complex numbers, and digging your own grave—are often typically thought of as more creative, at least until they become entrenched.) In two-sided networks, then, we expect to see increasing competition between optimality principles and increasingly many opportunities for failure to satisfy them.

The computer desktop provides an illustration of many of these competitions and opportunities for failure. First let us consider an aspect of the desktop blend in which Topology clashes with Integration, and Integration of the blend wins. The purpose of the blend is to provide an integrated conceptual space that can serve as the basis for integrated action. The basic integrative principle of the computer desktop is that everything is on the two-dimensional computer screen. But in the input space of real office work, the trashcan is not on the desktop. By Topology, the location of the trashcan as not on the desktop would be projected to the computer interface blend; but doing so would destroy the internal integration of the blend, which is why, on the computer screen desktop, the trashcan is on the desktop. Integration of the blend in this case can only be achieved by relaxing the topology constraint as we develop a new frame for the blend.

There are at least two reasons why we are content to relax topology in this way. First, the topology that is being dropped from the desktop input is incidental to the cross-space mapping—the three-dimensionality of the office and the position of trashcans under desks has no counterpart in the cross-space mapping to the input of computer operation. Second, as we have mentioned, the purpose of constructing this blend is to develop a conceptual basis for extended action, and not to draw conclusions about the input space of offices. In a contrasting case, like the Buddhist monk, the purpose is to draw conclusions about topology of input spaces—specifically coincidence of locations and times. In such a case, relaxing Topology is likely to allow inferences in the blend that would project wrongly or not at all back to the input, and so defeat the purpose of the blend. In that case, Topology is not relaxed.

It is also possible for the frame elaborated for a blend to fail to satisfy the optimality constraints. The most noticeable such failure for the computer desktop is the use of the trashcan both as the container of what is to be deleted and as the instrument of ejecting floppy disks. This failure involves failures of Integration, Topology, and Web.

The trashcan-for-both-deletion-and-ejection violates Integration for the frame elaborated for the blend in three ways. First, in the frame elaborated for the blend, the dual roles of the trashcan are contradictory, since one ejects the floppy disk to keep it rather than discard it. Second, in the frame elaborated for the blend, all other operations of dragging one icon to another have as their result that the first is *contained* in the second, but that is not so in the uniquely exceptional case of dragging the floppy to the trashcan. Third, for all other manipulations of icons on the desktop, the result is a *computation*, but in this case it is a physical *interaction* at the level of hardware.

The trashcan-for-both-deletion-and-ejection violates Topology. In the input of office spaces, putting an

object in a folder or in the trashcan results in containment. This topology is projected to the blend. The trashcan in the desktop is like any icon that represents a metaphoric container: if we drag a file to a folder icon or to the trashcan icon, the file is then deposited there, and this is the topology of the input of office spaces. However, putting the floppy disk icon into the trashcan icon so as to eject it is an exceptional and contrary case that violates the projection of topology from the input of offices. It also violates topology by not preserving the relation Input 2 (the space of real offices) that items transferred to the trashcan are unwanted and destined to become non-retrievable.

The trashcan-for-both-deletion-and-ejection also violates Web. The very opportunity of ejecting floppy disks from the computer desktop creates non-optimal web connections, since sometimes the floppy disk is "inside" the world of computer operations and sometimes it is "inside" the world of the real office.

We now turn to questions of optimality in word-processing programs on the desktop. The command sequence Select-Copy-Paste on word-processing applications violates both Topology and Web. It violates Topology as follows. In the Input where text is actually copied by scribes or Xerox machines, copying (after selection) is a one-step operation. There is no pasting and no clipboard. Properties specific to the Integration in the blend make it convenient to decompose this operation into two steps, but they do not map topologically onto corresponding operations in the Input of "real copying."

The labels "Copy" and "Paste" chosen for these two operations in the blend also violate Web: the Copy operation in the blend (which actually produces no visible change in the text) does not correspond to the Copy operation in the Input (which does produce visible change); the Paste operation, which does produce change, is closer to "copying" in the Input, but the label "Paste" suggests a counterpart (pasting), which is not even part of the copying process. Not surprisingly, these flaws in the overall blend lead to mistakes by novice users. They click Copy instead of Paste, or try sequences like: Select—Select Insertion Point—Copy. This fails miserably because the first selection (not marked for copying) is lost when the second selection occurs, and anyway Copy at that point is the wrong command. Mistakes like this are interesting however, because they represent the user's effort to maintain optimal Topology and Web connections. If double selection were possible on the blended interface (as it is, in terms of attention, in the Input), Copy and Paste could easily be reintegrated into a single process operating on both selections, and the attempted sequence would be viable. In fact, the Microsoft Word® application being used to type the present text has a keyboard command (with no counterpart in the menus) which comes closer to this conception.

The "Cut and Paste" method of moving text is a less severe violation, because the projected operations from the "office" Input are plausible and properly web-connected. But it does add conceptual complexity to what is more easily conceived of as simple unitary "moving." Recent versions of Word® have added to the interface the possibility of selecting and dragging text directly to the appropriate location. The portion of text does not actually "move" (only the arrow does) until the mouse is unclicked.

Despite all these failures to satisfy optimality principles, nonetheless the desktop blend draws rich and effective structure from familiar frames, and users are able to use it in a rudimentary fashion very quickly and

to learn the elaborated frame, warts and all. The non-optimality creates difficulty for novices, who are reluctant to put the floppy disk in the trashcan since by topology it should then be lost, but this difficulty is forgotten by advanced users, who learn a less optimal but more elaborate blend.

The fact that in two-sided networks the organizing frame of the blend is not available by extension from the organizing frame of either input increases chances of non-optimality and of competition between the optimality principles, but it also offers opportunity for creativity in attempting to satisfy the optimality principles. Pressure to satisfy optimality principles in highly complex two-sided shared topology networks has historically given rise to some of the most fundamental and creative scientific discoveries. The development of the concept of complex numbers in mathematics, discussed in chapter four, is a case in point. The complex number blend turns out to be a two-sided shared topology network. Some key elements in each input have no counterparts in the basic cross-space mapping. The operation of multiplication for numbers has no counterpart in the geometry input, and the angles of vectors in the geometry input have no counterparts in the number input. The blend, however, inherits both the multiplication operation from the frame of the "number" input and the vector angle from the frame of the "geometric" input. This is already enough to make it a two-sided shared topology network, since multiplication in the blend has TF2 topology while angle in the blend has TF1 topology. But furthermore, in the blend, multiplication includes addition of angles as one of its constitutive components. This is discovered by running the blend; it turns out to be a highly unexpected essential property of the new concept of number which has emerged. So in this instance, the pressures to satisfy optimality in this two-sided shared topology network led to important mathematical discovery. Jeff Lansing has pointed out to us other marvelous examples of important scientific blends leading to discovery (by Fourier, Maxwell, and Faraday), which suggests that this is a general process. We emphasize that this type of creativity is possible by virtue of the competition of optimality principles and the power of blending to accommodate them.

Unpacking is actually relatively easy to satisfy in the two-sided case since key elements in the blend cannot all be projected back to the same organizing frame of one of the inputs. For example, in *digging your own grave*, the gravedigger is responsible for the death, and this structure cannot be provided by the single organizing frame of digging graves, making it clear that the blend must be unpacked to the organizing frames of different inputs.

Competition among pressures motivated by optimality

At the top level of our model, there are general structural characteristics of all blend structures, like cross-space mappings. At a lower level, there are optimality principles like Integration. But these optimality principles themselves compete, as we have seen and as we will discuss further, and that competition results in a variety of yet lower-level optimality pressures for constructing the blend. In this section, we discuss the candidate optimality pressures for which we have found evidence.

Non-disintegration: Neutralize projections and topological relations that would dis-integrate the blend.

For example, as we saw above in the section called "Optimality principles and two-sided networks," since the integrative principle of the computer desktop blend is that everything is on the computer desktop, Topology must be relaxed in projecting the trashcan to the blend so as to filter out the three-dimensionality of the real office space. In the regatta example, weather in 1853 (even if known) is not projected because it would clash with the projected 1993 weather. In the Debate with Kant, the language, German, from Input 1, is not projected; Integration in the debate frame requires a single language.

Non-displacement: Do not disconnect valuable web connections to inputs.

The computer desktop has web connections to the space of computer operations, in which all shifts of focus require only a simple click. For example, if a user is running five different applications on the desktop and wants to see only one of them, he can click "Hide Others" (conversely, "Show others"); to see a given document partially occluded by another, he need only click anywhere on the desired document. But in the space of offices, to hide everything on your desk except the one thing you wish to focus on would require complex physical operations. If these operations were all projected to the blend, it would sever its useful web connections to the input of computer operation. Function guides competition here. The web connection to "change of focus" in the computer operations input is important because the desktop interface is designed to run a computer. If its function were to simulate the working environment of an office worker, then the complexity of the physical operations would be maintained at the expense of computing efficiency.

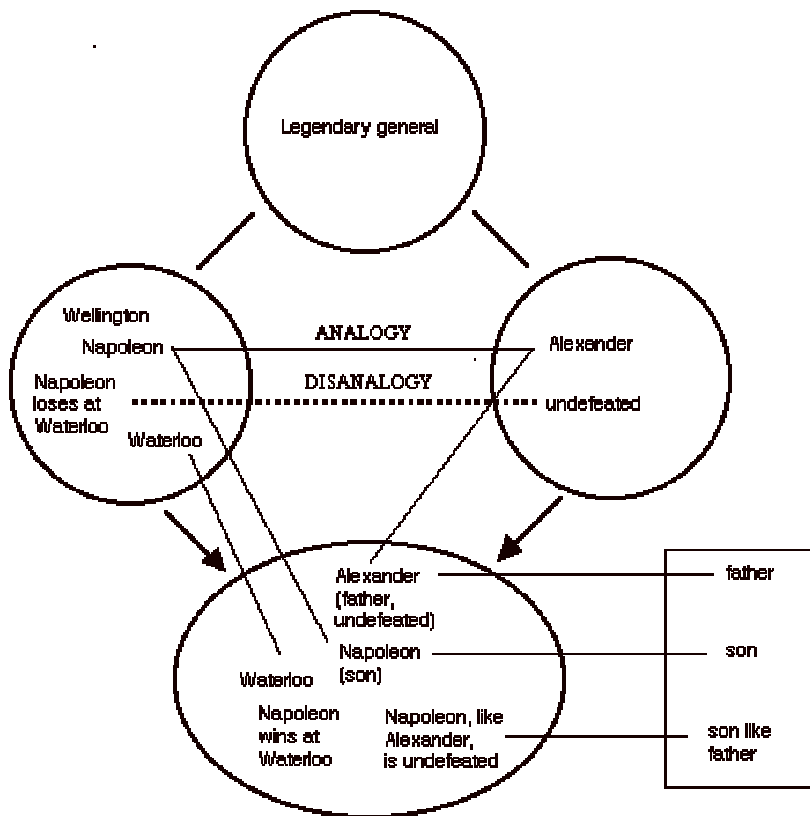
Non-displacement combines with Integration to force novel integrations in the blend. For instance, in the case of the metaphor "digging one's own grave," the blend's causal, temporal, and intentional structure (*digger is unaware of his actions, a deep enough grave causes death*) are projected from the target space of mistakes and failure. This web connection is crucial to the reasoning, but would be destroyed if the commonplace structure of the source (death followed by conscious grave-digging by somebody else) were projected. In the Nixon-in-France example, we project to the blend Nixon, but not his U.S. citizenship, which would prevent him from being president of France, thus cutting off a crucial web link from the blend to the second input.

Non-interference: Avoid projections from input spaces to the blend that defeat each other in the blend.

For example, in the space of office work, we often write "discard" across the top of an outdated version to be discarded. In the computer desktop, the icon for a file has only one place for a label. If we projected to the computer desktop the operation of labeling the document "discard" by making a click-command that put that label on the icon, we would lose the title of the file. So the "title" label and the "discard" label from the space of office work defeat each other if both are projected to the blend.

In a counterfactual blend like "If Napoleon had been the son of Alexander, he would have won the battle of Waterloo," we do not attempt to project Napoleon's actual father. There is no inference that if Napoleon had been the son of Alexander, Charles Bonaparte would have been Alexander, although formally this leads to an

integrated scene. (If the goal is to point out that Napoleon lacked some military qualities that perhaps Alexander possessed, it is odd to say "If Charles Bonaparte had been Alexander, Napoleon would have won the battle of Waterloo." But if we really mean that some deficiency was actually transmitted by Charles, through genes or education, then the counterfactual sounds ok. The blend in that case contains an efficient father appropriately connected to Charles and Alexander.) The traits of fathers Alexander and Charles would defeat each other in the blend.



In the metaphor of Death as the Grim Reaper, Death, which inherits from the cadaver its skeleton, could additionally inherit its shredded, decayed clothing. The cowl would thus be shredded, but this interferes with the projection of the cowl as a piece of clothing of a live priest at the funeral. Technically, the projections from the blend to the input, of the skeleton, the sickle, etc., are one-to-one, but the projection of the shredded cowl would be one-to-many (the priest's head-dress/the dead man's hat). Similarly, in a "bad" desktop blend, the projection (from the blend to the space of real offices) of the label on the desktop file would be one-to-many: the title of a document, or the instruction to discard it.

Non-ambiguity: Do not create ambiguity in the blend that interferes with the computation.

The method of ejecting floppy disks "through" the trash on the computer desktop violates several constraints, as we have seen above. It also violates non-ambiguity. Superposition of icon a over icon b "means" copying/inserting the contents of a "into" b. So a plausible interpretation of the disk icon's being moved over the trashcan icon is that the contents of the disk are transferred to the trashcan. But in fact, the meaning in this particular case is entirely different ("eject disk from computer"). This makes the superposition schema in the blend ambiguous. Similarly, a debate-blend, which works with Kant and philosophy, might fail with a deity or prophet and religion, because we would not know whether to count victory in the debate as superior religious insight or as heresy. We would not have a straightforward way of running the ambiguous blend.

A cartoon blend advertising the magazine *Success* has a man blended into a rocket shooting into outer space. People judge this to be a "bad metaphor." One reason, presumably, is the inherent ambiguity in the blend: it is good for a rocket to fly, but not good for a man to be shot out of a cannon with no control over his actions and fate.

Backward projection: As the blend is run and develops emergent structure, avoid backward projection to an input that will disrupt the integration of the input itself.

During blending, conceptual work may be performed at any site in the conceptual array. For example, one straightforward way to optimize Topology is to project the topology of the blend back to reform the inputs. But doing so will conflict with the original Integration of the inputs. Usually, this is undesirable, which gives rise to pressure to avoid backward projection.

For example, under pressure from Integration, the desktop blend places the trashcan on the desktop, but projecting this relation backward to the input of the actual desks would disrupt their efficient use. In the grave-digging metaphor, we do not want to start thinking, through backward projection to the source input, that digging graves actually causes death. We do not interpret the printing press cartoon as additionally suggesting that smashing cars with a printing press is a good idea.

Many blends, however, have the purpose of modifying the structure of an input. Coulson (1995) considers such blends.

The Topology Constraint and the Invariance Hypothesis

One goal of the network model is to account for inferencing during conceptual projection. For example, we have shown in our pedagogical riddle of the Buddhist monk that if the blend and its inputs have the same co-occurrence of locations and times (under Topology) and this mirroring survives as we run the blend (under

Web), then the inference of an *encounter* in the blend entails inferences for the inputs which effectively solve the riddle.

An earlier attempt to account for inferencing during conceptual projection in the special case of metaphor goes under the name of "the invariance principle"—launched by analysis in Turner (1987: 143-148), stated briefly in Lakoff and Turner (1989: 82), and analyzed in Lakoff (1989), Turner (1991: 172-182), Lakoff (1993), and Turner (1996b). The invariance principle proposes that in metaphor, we attempt to project image-schematic structure (with inferences) from source to target while avoiding the creation of an image-schematic clash in the target. Importing new image-schematic structure to the target by projection does not violate the invariance principle if the original target is appropriately indeterminate. Asserting by means of the metaphor that the target's image-schematic structure is to be overridden does not violate the constraint, since the changed target contains no clash.

The network model of conceptual projection extends and modifies the invariance principle. We emphasize the importance of image-schematic topology in all conceptual projection, not only metaphoric projection. In the network model, there are productive matches of image-schematic structure between inputs, generic space, and blend. First, consider the generic space and the inputs. The structure of the finished generic space, taken as applying to both inputs, frequently contains extensive image-schematic structure, as in the riddle of the Buddhist monk, where the two input spaces do not stand in metaphoric relation (we do not understand the descending monk by metaphoric projection from the ascending monk or conversely).

Second, consider the blend and the inputs. There is always important matching of image-schematic topology between blend and inputs under Topology: the Buddhist monk blend requires an extensive topological match between parts of the blend and each of the inputs.

But the Topology constraint is not a generalization of the invariance principle to non-metaphoric cases. The Topology constraint does not require that we project image-schematic structure from one input to the other or from the blend to the inputs. It does require that we project image-schematic structure from the inputs to the blend. In the Buddhist monk riddle, we do not import image-schematic structure from one input to the other, because *the detailed relevant image-schematic structure already exists in each input independently of the other input*. Furthermore, although we "understand" the Buddhist monk input spaces by drawing on the image-schematic structure of the "encounter" in the blend, we do not project that image-schematic structure from the blend to the inputs; quite the contrary. The blend has the image-schema for "encounter"; the inputs do not have it; we do not project that image-schema to the inputs; instead, we infer from this image-schema in the blend a different and complicated relation of image-schemas between the inputs: namely, there exists in each input a time-location pair, and these two pairs in the two inputs have the identical times and the identical locations.

Topological structure in the blend may be elaborated that is important for the construction of meaning but that is not projected identically back to the inputs. This is clearest in the case of science fiction or fantasy blends meant for entertainment, where we are not solving over the inputs, and where Topology and Web may

be thoroughly relaxed, but it is also true for cases where inferences are drawn for the inputs: the existence of the race in the *Great America II* blend is crucial for the construction of meaning and inference, but the race structure in the blend does not displace the structure of the inputs in which each boat is making a solitary run. Each space in the conceptual projection has a different structure, and each space is useful.

Now let us consider examples that are felt to be clearly metaphoric. What is the relationship for clearly metaphoric cases between the topology constraint—which we claim applies to all integration networks—and the invariance principle—which was advanced exclusively for the metaphoric cases? The network model, far from eliminating the need for a theory of metaphor and a consideration of the mapping of image-schemas, requires such a theory, in the following way. Consider the status of the generic space and the origin of its content. Typically, the generic space contains image-schematic topology which is taken to apply to two inputs. Often, much or even all of that content is supplied processually by activating a conventional metaphoric mapping between the domains underlying the two inputs. Indeed, in many cases, some of them quite important, it may be that the image-schematic structure belongs to the target only because metaphoric projection installed it in the target. In sum, a counterpart mapping is needed to launch on-line blending, and that counterpart structure is often supplied by activating a conventional metaphor, and the counterpart structure may have been created by the basic metaphor projection rather than merely picked out as a template for the projection.

Now consider the case where the metaphoric meaning that arises in an integration network is not supplied by activating a conventional conceptual metaphor. In these cases, the invariance principle survives with modification into our model. Under the topology and web constraints, the projection of image-schematic structure from the source space plays an important role in blending. Under the topology constraint and the non-disintegration pressure for the inputs, image-schematic clashes are avoided in the target space. Moreover, if there is a clash of topology between source and target, then since it is the target we care about, we typically prefer the topology in the target: structure in the blend needed to deliver inferences for the target will accord with the important image-schematic structure in the target as opposed to the source. We see a clear example of this in *digging your own grave*, where the causal, intentional, frame, and internal event structure of the blend suit the topology of the target space but not at all that of the source space, although some structure—the foolishness of failing to recognize concrete actions—comes from the source into the blend. In general, the topology of the blend needed for delivering inferences for the target cannot do so if it conflicts with the protected topology of the target. A clash of this nature is to be avoided. This principle is equivalent in spirit and effect to the invariance principle's proposal that an image-schematic clash is not to be created in the target.

But the network model and its Topology principle differ from the two-domain model of metaphor and its invariance principle. Under the invariance principle, all the inferential structure had to be supplied by either the target and its protected image-schematic structure or by the source image-schematic structure projected to the target. We hope we have demonstrated that the blend often has emergent structure available from neither input but important for inferencing. In *digging your own grave*, there is important causal structure and event structure: the person addressed is digging a grave and the existence of a satisfactory grave causes death. This structure is image-schematic, but it is not given by either input. The causal structure of the blend is the inverse of the causal structure of the source, and in the target it is not given, prior to the blend, that the person

addressed is performing bad acts, that performing them completes in a cumulative manner a certain gradual action, or that completing that action causes disaster. This image-schematic structure, with its inferences, is developed in the blend so as to permit the projection of certain inferences to the target that the target can accept.

Similarly, the desktop has emergent structure provided by neither input, such as dragging a file icon from the hard disk icon to the floppy disk icon to *duplicate* the file onto the floppy disk rather than *move* it off the hard disk and onto the floppy disk. The image-schematic topology of the blend in this instance violates the topology of the source of actual desktops and moving things on them, and it is not given by the target space of symbolic computer commands, although it can be projected there.

There is another important difference. The two-domain model of metaphor with its invariance principle is not a theory of the development of metaphoric mappings. In our view, the development of a conventional metaphoric mapping involves conceptual integration. In cases where useful inferences or structure have been projected from the blend to the target so that the mapping from source to target becomes thoroughly conventional, and the blend is no longer a working space, it is possible to overlook both blend and generic space. Additionally, blending is always available to someone who activates a conventional metaphor, and many of the conventional metaphors studied so far are, like *ANGER IS HEAT* or *The Grim Reaper*, actually conventional blends.

VII. Additional dimensions of conceptual integration

Activation

In Fauconnier & Turner (1994), we provided a taxonomy of blends by kind of conceptual activation. The parameters in this taxonomy are: the number and type of spaces involved; the degree to which any particular space in the array is active as a working space in which new on-line conceptual construction must be done; the degree of blending and of abstraction; whether the vocabulary transfer is on-line or permanent; the number of conceptual domains involved in building up the inputs and the blend; whether or not the conceptual domain involved is consciously focused upon; and the extent to which the blended space gives birth to a new conceptual domain. The existence of a blended space does not entail that it serve as the basis for an imaginary conceptual domain, like the ghost ships of the boatrace example or the sinners of Dante's hell. Most blends, while serving important local cognitive functions, have no corresponding conceptual domains.

Functions of blends and topic spaces

The function of the desktop blend is to provide an integrated activity that the computer user can inhabit; naturally, the integration principle dominates. But in the monk example, the function of the blend is to solve

for a puzzle in the inputs; naturally, the web principle dominates. The many examples analyzed in previous work on blending supply a survey of functions of blends. They include: reasoning on inputs (the monk example); adding meaning and emotion to inputs (enthusiasm in the boat race example); creating rhetorical presence (Oakley 1995) for some aspect of the inputs ("If gnatcatchers were dolphins, we would not be permitting them to become extinct"); jokes (analyzed by Seana Coulson); conceptual change ("artificial life"); cultural change ("same-sex marriage"); provisional category extension ("he's a real fish"); enhancing one of the inputs (the debate with Kant enhances the modern philosophy teacher's authority, status, etc.); supplying new action (desktop); providing integrated conceptual structure over an unintegrated array (as in giving the structure of caused motion to unintegrated events in "John sped the toy car around the Christmas tree"); and integrating the performance of actions (learning to ski). It is important to remember that functions cannot be predicted from structural features.

For functional reasons, the input spaces are rhetorically unequal. For example, in the boat race, it is 1993 that the reporter cares about and talks about. It is 1993 that he is interested in understanding and reporting fully. We label 1993 the *topic space* of the projection. It is possible for there to be more than one topic space (in the monk example, both spaces are equally topic spaces). It is also possible for the topic space to shift: if we are descendants of the captain of *Northern Light*, it may be 1853 that we care about understanding. Coulson (1995) shows that a source input space in a metaphoric projection can be the topic space.

VIII. Summary and further results

Conceptual integration—"blending"—is a basic cognitive operation. Integration networks involve input spaces, generic spaces, and blended spaces. There is a cross-space mapping of counterpart connections between input spaces and selective projection of structure from the inputs to the blend. Blends develop through composition, completion, and elaboration. Blends provide the possibility of backward projection to the inputs of inferential and other structure. Conceptual integration networks arise under competing optimality principles of Topology, Integration, Web, Unpacking, Good Reason, and Metonymy Projection. Some basic patterns of satisfying these constraints are single-framing networks, frame networks, and shared topology networks. Shared topology networks are one-sided or two-sided. Two-sided shared topology networks can be asymmetric or unfilled, and are further distinguished by variation in the projection of secondary topology from an input to a blend, below the level of organizing frame.

There are many other results of this research that can only be referred to here, without further explanation. We provide an analysis of grammatical constructions used to evoke conceptual integration, and of the way those grammatical constructions can be composed to evoke compositions of conceptual integrations. We analyze the mechanisms of frame integration, including composition of frame integrations. Unsurprisingly, we find that the construction of meaning is not truth-conditionally compositional: construction of meaning is not just a matter of specifying contextual elements and composing truth conditions. However, it turns out that there *is* compositionality at the level of the general schemes for conceptual integration networks and at the level of the syntactic forms that prompt for those schemes. We show that one purpose of grammatical constructions is to prompt for conceptual integrations of certain types. We also show that there is a process of *formal* blending at the level of grammar that is parallel to the process of *conceptual* blending, and that the

two processes interact in intricate ways. In particular, conceptual blending can guide formal blending to produce new grammatical constructions suited to evoke just those conceptual blends. In these ways, blending is a central process of grammar. We analyze the role of conceptual integration in conceiving of space, form, and motion. We explore typical uses of conceptual integration in literature and the visual arts. We argue that conceptual integration interacts with cognitive activities like category assignment, analogy, metaphor, framing, metonymy, grammatical constructions, and so on. Moreover, the model of conceptual intergation suggests that these are not sharply distinguished kinds of cognitive activity. Distinctions among the products of conceptual integration are real, but arise from a number of interacting graded dimensions of difference. A number of locations in the grid of all these dimensions stand out as prototypes or cognitive reference points, and these locations have been given the name "categorization," "framing," "metaphor," and so on. These reference points are convenient, but there are not sharp divisions in the very nature of the types of phenomena that fall under these labels. The underlying cognitive operations are general, while the differences stem from the nature of the appropriate domains and mappings and the many interacting dimensions along which they vary.

These lines of research have also been pursued in some detail by Coulson (1997), Mandelblit (1997), Oakley (1995), Robert (in press), Sweetser (to appear), Freeman (1997), Grush and Mandelblit (ms), Veale (ms), Zbikowski (ms).

IX. Conclusion

This paper has presented evidence for a general cognitive operation, conceptual integration, which builds up networks of connected spaces—inputs, generic, and blended spaces. The construction of such networks depends quite generally on establishing cross-space mappings of the sort commonly studied in theories of metaphor and analogy. But metaphor and analogy phenomena are only a subset of the range of conceptual integration phenomena. Conceptual integration networks are equally prominent in counterfactuals, category extension, event integration, grammatical constructions, conceptual change (as in scientific evolution), and literary and rhetorical invention. The salient feature of such networks is the construction of a blended space which develops specific emergent structure and dynamics while remaining linked to the overall network. Projection in a network can occur in different ways and in different directions, as we analyze in the taxonomy of section VI. Theories of metaphor and analogy have typically focused on the case where projection is one-way (from a "source" to a "target") and they have overlooked the construction of blended spaces. Accordingly, the overall picture is even richer than previously envisioned, and any explicit computational modeling of the entire process will presumably face obstacles in addition to the already formidable ones encountered for analogy. At the same time, however, the conceptual integration view yields a far more unified general conception of meaning construction at all levels, and this should prove to be a major simplification.

It is remarkable that blending—a general-purpose, fundamental, indispensable cognitive operation, routinely employed in a variety of domains, commonly interactive with other cognitive operations that have received extensive analysis—should have received so little systematic attention in the study of cognition and language. The routine and largely unconscious nature of blending may have helped it escape scrutiny. The

many well-known spectacular blends—sirens, mermaids, chimerae, space aliens, cybernetic organisms, Bambi—may have made blending seem merely exotic. We hope to have demonstrated that blending is a central, orderly, powerful, systematic, and commonplace cognitive operation. We have proposed a theoretical model of its structural and dynamic principles.

References

Barsalou, L. 1996. *Perceptual Symbol Systems*. University of Chicago ms.

Burns, Bruce. 1995. *Fluid Concepts and Creative Analogy: A Review*. *AI Magazine*, 16.3. pp. 81-83.

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

Coulson, Seana. 1997. *Semantic Leaps: Frame-Shifting and Conceptual Blending*. UCSD Ph. D. dissertation.

Cutrer, Michelle. 1994. *Time and Tense in Narratives and Everyday Language*. Doctoral dissertation, University of California at San Diego.

Dinsmore, J. 1991. *Partitioned Representations*. Dordrecht: Kluwer.

Fauconnier, G. 1994. *Mental Spaces*. New York: Cambridge University Press. [Originally published (1985) Cambridge: MIT Press.]

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

Fauconnier, G. 1990. Domains and Connections. *Cognitive Linguistics* 1.1.

Fauconnier, G. & E. Sweetser. 1996. *Spaces, Worlds, and Grammar*. Chicago: University of Chicago Press.

Fauconnier, G. & Turner, M. 1994. *Conceptual projection and middle spaces*. UCSD: Department of Cognitive Science Technical Report 9401. [Available from <http://cogsci.ucsd.edu> and from <http://markturner.org>]

Fauconnier, G. & Turner, M. 1996. Blending as a central process of grammar. In *Conceptual Structure, Discourse, and Language*, Ed. Adele Goldberg. Stanford: Center for the Study of Language and Information [distributed by Cambridge University Press].

Fauconnier, G. & Turner, M. (in preparation). *Conceptual Integration*.

Fillmore, C. and P. Kay. n.d. ms. On Grammatical Constructions. University of California at Berkeley.

Fong, H. 1988. *The Stony Idiom of the Brain: A Study in the Syntax and Semantics of Metaphors*. UCSD Ph.D. dissertation.

Forbus, Kenneth D., Dedre Gentner, and Keith Law. 1994. "MAC/FAC: A Model of Similarity-based Retrieval." *Cognitive Science* 19, 141-205.

Forbus, Kenneth D., Dedre Gentner, Arthur Markman, and Ronald Ferguson. 1997. Analogy just looks like high level perception: Why a domain-general approach to analogical mapping is right. Manuscript, Northwestern University.

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

French, R. M. 1995. *The subtlety of sameness*. Cambridge, MA: The MIT Press.

Gentner, Dedre. 1983. "Structure-Mapping: A Theoretical Framework for Analogy." *Cognitive Science* 7, 155-170.

Gentner, Dedre. 1989. "The mechanisms of analogical reasoning. In S. Vosniadou and A Ortony (Eds.), *Similarity and analogical reasoning*. Cambridge: Cambridge University Press.

Gibbs, R. W., Jr. 1994. *The Poetics of Mind: Figurative Thought, Language, and Understanding*. Cambridge: Cambridge University Press.

Glucksberg, S. & Keysar, B. 1990. Understanding metaphorical comparisons: Beyond similarity. *Psychological Review* 97, 3-18.

- Goffman, E. 1974. *Frame Analysis*. New York: Harper and Row.
- Goldberg, A. 1994. *Constructions*. Chicago: University of Chicago Press.
- Goodman, N. 1947. The problem of counterfactual conditionals. *Journal of Philosophy* 44.
- Grush, R. 1995. Cognition and Emulation. Ph.D. dissertation, University of California San Diego.
- Hofstadter, D. 1985. Analogies and Roles in Human and Machine Thinking. *Metamagical Themas*. Chap. 24. New York: Bantam Books.
- Hofstadter, D. 1995a. *Fluid Concepts and Creative Analogies*. New York: Basic Books.
- Hofstadter, D. 1995b. A Review of *Mental Leaps: Analogy in Creative Thought*. *AI Magazine*, 16.3. pp. 75-80.
- Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. 1986. *Induction: Processes of inference, learning, and discovery*. Cambridge, MA: Bradford Books/MIT Press.
- Holyoak, Keith J. and Paul Thagard. 1989. "Analogical Mapping by Constraint Satisfaction." *Cognitive Science* 13, 295-355.
- Holyoak, Keith J. and Paul Thagard. 1994. *Mental Leaps: Analogy in Creative Thought*. MIT Press: Cambridge, Mass.
- Hummel, John. and Keith Holyoak. 1996. LISA: A Computational Model of Analogical Inference and Schema Induction. *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society*. Edited by Garrison Cottrell. Lawrence Erlbaum: Mahwah, New Jersey.
- Hutchins, E. 1995. *Cognition in the Wild*. Cambridge: MIT Press.
- Indurkha, Bipin. 1992. *Metaphor and Cognition: An interactionist approach*. Dordrecht: Kluwer.
- Johnson-Laird, P.N. 1983. *Mental Models*. Cambridge, Mass.: Harvard University Press.

Kahneman, Daniel. 1995. Varieties of counterfactual thinking. In *What Might Have Been: The Social Psychology of Counterfactual Thinking*. Edited by Neal J. Roese and James M. Olson. Lawrence Erlbaum Associates: Hillsdale, N.J.

Keane, Mark T., Time Ledgeway, and Stuart Duff. 1994. "Constraints on Analogical Mapping: A Comparison of Three Models." *Cognitive Science* 18, 387-438.

Kline, M. (1980). *Mathematics. The Loss of Certainty*. Oxford: Oxford University Press.

Koestler, Arthur. 1964. *The Act of Creation*. NY: Macmillan.

Kunda, Z., D. T. Miller, and T. Clare. 1990. Combining social concepts: the role of causal reasoning. *Cognitive Science* 14. 551-577.

Lakoff, George. 1987. *Women, Fire, and Dangerous Things* (Chicago: University of Chicago Press), case study 1, pages 380-415.

Lakoff, George. 1989. The Invariance Hypothesis: Do Metaphors Preserve Cognitive Topology? Paper delivered at the 1989 Duisburg Cognitive Linguistics Conference and available from the Linguistic Agency at the University of Duisburg).

Lakoff, George. 1993. The Contemporary Theory of Metaphor. In Andrew Ortony, editor, *Metaphor and Thought*, second edition (Cambridge: Cambridge University Press) pages 202-251.

Lakoff, George and Mark Johnson. 1980. *Metaphors we live by*. Chicago: University of Chicago Press.

[LAKOFF & NUNEZ - INCLUDE CITATION HERE.]

Lakoff, George and Mark Turner. 1989. *More than Cool Reason: A Field Guide to Poetic Metaphor* (Chicago: University of Chicago Press).

Langacker, R. 1987/1991. *Foundations of Cognitive Grammar. Vol.I, II*, Stanford University Press.

Langacker, R. 1993. Reference-point constructions. *Cognitive Linguistics* 4.1: pp. 1-38.

Lewis, David. 1972. *We, the Navigators. The Ancient Art of Landfinding in the Pacific*. Honolulu: The University Press of Hawaii.

Mandelblit, Nili and Oron Zachar. (in press). The Notion of Unit and its Development in Cognitive Science. *Cognitive Science*.

Mandelblit, Nili 1997. *Grammatical Blending: Creative and Schematic*

Aspects in Sentence Processing and Translation. Ph.D. dissertation, UC San Diego.

Mitchell, M. 1993. *Analogy-making as perception*. Cambridge: MIT Press.

Moser, D. and D. Hofstadter. (undated ms.) Errors: A Royal Road to the Mind. *Center for Research on Concepts and Cognition*. Indiana University.

Nunberg, G. 1978. *The Pragmatics of Reference*. Bloomington, Ind.: Indiana University Linguistics Club.

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

Ortony, A. 1979a. Beyond literal similarity. *Psychological Review* 86, 161-180.

Ortony, A. 1979b. The role of similarity in similes and metaphors. In A. Ortony, editor, *Metaphor and thought*. Cambridge: Cambridge University Press, pages 186-201.

Robert, Adrian. (in press) "Blending in the interpretation of mathematical proofs." In: *Conceptual Structure, Discourse, and Language, II* Edited by Jean-Pierre Koenig, et al. Stanford: Center for the Study of Language and Information (CSLI) [distributed by Cambridge University Press].

Schwartz, Daniel L. and John B. Black. 1996. "Shuttling Between Depictive Models and Abstract Rules: Induction and Fallback." *Cognitive Science* 20, 457-497.

Sweetser, E. 1990. *From etymology to pragmatics: the mind-as-body metaphor in semantic structure and semantic change*. Cambridge: Cambridge University Press.

- Sweetser, Eve. 1997. "Mental Spaces and Cognitive Linguistics: A Cognitively Realistic Approach to Compositionality." Fifth International Cognitive Linguistics Conference. 1997.
- Talmy, L. 1977. "Rubber-sheet Cognition in Language." *Proceedings of the 13th Regional Meeting of the Chicago Linguistic Society*.
- Talmy, L. 1978. "Figure and Ground in Complex Sentences." In: *Universals of Human Language: Syntax* (vol. 4). Edited by Joseph Greenberg. Stanford University Press.
- Tetlock, Philip and Aaron Belkin. 1996. Counterfactual Thought Experiments in World Politics: Logical, Methodological, and Psychological Perspectives. In Philip E. Tetlock and Aaron Belkin, editors, *Counterfactual Thought Experiments in World Politics: Logical, Methodological, and Psychological Perspectives*. Princeton: Princeton University Press, 1996, pages 3-38.
- Turner, Mark. 1987. *Death is the Mother of Beauty: Mind, Metaphor, Criticism* (Chicago: University of Chicago Press).
- Turner, Mark. 1991. *Reading Minds: The Study of English in the Age of Cognitive Science*. Princeton: Princeton University Press.
- Turner, Mark. 1996a. Conceptual Blending and Counterfactual Argument in the Social and Behavioral Sciences in Philip Tetlock and Aaron Belkin, editors, *Counterfactual Thought Experiments in World Politics* (Princeton: Princeton University Press).
- Turner, Mark. 1996b. *The Literary Mind*. (New York: Oxford University Press).
- Turner, Mark and Gilles Fauconnier (1995). Conceptual Integration and Formal Expression. *Journal of Metaphor and Symbolic Activity*, volume 10, number 3, pages 183-204.
- Turner, Mark and Gilles Fauconnier (in press). Conceptual Integration in Counterfactuals. In *Conceptual Structure, Discourse, and Language, II*, Ed. J.-P. Koenig. Stanford: Center for the Study of Language and Information [distributed by Cambridge University Press].
- Van Hoek, Karen. 1997. *Anaphora and Conceptual Structure*. Chicago: The University of Chicago Press.

Veale, Tony. 1996. Pastiche: A Metaphor-centred Computational Model of Conceptual Blending, with special reference to Cinematic Borrowing. Manuscript.

Zbikowski, Lawrence. 1997. "Conceptual blending and song." Manuscript, University of Chicago.

Zribi-Hertz, Anne. 1989. Anaphor Binding and Narrative Point of View: English Reflexive Pronouns in Sentence and Discourse. *Language* 65, pp. 695-727.

Footnotes

[1] There is widespread agreement in research on analogy and metaphor that cross-space mappings operate and transfer inferences by extracting or creating common schematic structure. The modeling of such processes has typically focused on the stage at which two domains are already appropriately structured and alignment takes place. Most researchers acknowledge, however, that this is only a part (perhaps even a small part) of the entire process, given the richness of domains and the corresponding multitude of ways to structure them (or "re-represent" them). These issues are discussed in many places (e.g. Burns (1995), Hofstadter (1995), Hofstadter (1995a, 1995b), Holyoak and Thagard (1994), Forbus et al. (1997), Hummel and Holyoak (1996)). The work we present in this article does not bear directly on this issue (but see footnote 3). It takes as given the undeniable, but admittedly still poorly understood cognitive capacity for schema induction and cross-domain mapping.

What we find with respect to cross-space mappings is:

- they operate in many phenomena other than metaphor and analogy;
- they operate extensively in the construction of simple everyday sentence meaning;
- they operate not just between a source and a target, but more generally between the various spaces of a conceptual integration network, including generic and blended spaces.

Our analyses of conceptual integration do, inevitably, have some consequences for the research on cross-space mapping. For example, we find evidence against all three of the claims in Dedre Gentner's classic paper on structure mapping (Gentner, 1983). (1) We find that, as a general principle, analogy is not compositional; the meaning of an analogy does not derive from the meaning of its parts. For example, "This surgeon is a butcher" has as part of its central meaning "incompetence," which is not available from either the input for the surgeon or the input for the butcher, but which is emergent in the blend. Personifying Death as a magician who is evil because he makes people disappear depends upon the emergence of *evil* in the blend: absent the blend, Death is not intentional and hence not evil, and a magician who performs disappearing tricks is not evil either. (2) We find, as a general principle, that mapping does depend upon specific content of the domains and not just on structural properties: the attribution of incompetence to the surgeon-butcher depends upon attitudes toward what happens to human bodies. (3) We find, as a general principle, that there

are not clean distinctions in kind between various products of conceptual projection and conceptual integration, but rather several interacting gradients of distinction. On the other hand, we concur in general with Holyoak & Thagard (1989) and Holland, Holyoak, Nisbett, & Thagard (1986) that pragmatic goals and purposes influence mapping, and with Keane, Ledgeway, and Duff (1994) that cognitive constraints (including, e.g., constraints on working memory, influence from background knowledge, influence of prior activity) influence mapping.

[2] For example, if $(r, \vartheta) = a + bi$ and

$(r', \vartheta') = a' + b'i$, then

$(r, \vartheta) \times (r', \vartheta') = (rr', \vartheta + \vartheta') =$

$(a+bi) \times (a'+b'i) = aa'-bb' + (a'b+ab')i$

[3] Douglas Hofstadter (personal communication) reports his discovery of how to "make" new geometries by blending. Taking projective geometry as a generic, and Euclidean as a source, he obtained a dual target for the latter, and a new "contrajective" geometry as a blend of the Euclidean and the Euclidal. Adrian Robert (in press) has shown that informal proofs in mathematics involve massive on-line blending of schematic structures, performed unconsciously by authors and readers of proofs.

[4] "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." (NY Times). This is an example of the Caused Motion construction studied in particular by Goldberg (1994), who explicitly addresses the issue of fusing grammatical constructions, within the framework of Construction Grammar (Fillmore and Kay 1999). We see this fusion as the reflex of conceptual blending. Fauconnier and Turner (1996) and Mandelblat (1997) offer detailed accounts of the Causative Construction in French and Hebrew respectively, using the blending approach. We also see Langacker's general approach to grammar as very congenial to the one described here. In Langacker's Cognitive Grammar, schemas are put in correspondence and integrated in succession to form functional assemblies. Interestingly, emergent structure also develops at this elementary level of sentence formation.