The Invariance Hypothesis: is abstract reason based on image-schemas?

GEORGE LAKOFF[1]

Abstract

I view cognitive linguistics as defined by the commitment to characterize the full range of linguistic generalizations while being faithful to empirical discoveries about the nature of the mind/brain.

The Invariance Hypothesis is a proposed general principle intended to characterize a broad range or regularities in both our conceptual and linguistic systems. Given that all metaphorical mappings are partial, the Invariance Hypothesis claims that the portion of the source domain structure that is mapped preserves cognitive topology (though, of course, not all the cognitive topology of the source domain need be mapped). Since the cognitive topology of image-schemas determines their inference patterns, the Invariance Hypothesis claims that imagistic reasoning patterns are mapped onto abstract reasoning patterns via metaphorical mappings. It entails that at least some (and perhaps all) abstract reasoning is a metaphorical version of image-based reasoning.

The data covered by the Invariance Hypothesis includes the metaphorical understanding of time, states, events, actions, purposes, means, causes, modalities, linear scales, and categories. Because the source domains of these metaphorical concepts are structured by image-schemas, the Invariance Hypothesis suggests that reasoning involving these concepts is fundamentally image-based. This includes the subject matter of Boolean, scalar, modal, temporal, and causal reasoning. These cases cover such a large range of abstract reasoning that the question naturally arises as to whether all abstract human reasoning is a metaphorical version of imagistic reasoning. I see this as a major question for future research in cognitive linguistics.

1. What is cognitive linguistics?

I generally prefer not to engage in methodological discussions and would rather just get on with my work. But I feel that the formation of a new

Cognitive Linguistics 1–1 (1990), 39–74 0936–5907/90/0001–0039 $2.00

journal devoted to cognitive linguistics calls for at least some discussion of what cognitive linguistics is, or at least what I take it to be. This is, of course, a personal statement. I include it because I would like to make the discussion of philosophical foundations and initial commitments a part of this enterprise from the outset.

It is my opinion that much of the acrimonious bickering that has characterized generative linguistics throughout its history has been due to a failure to engage in such discussions and to a lack of charity toward the primary commitments of others. I hope that if we make our primary commitments clear to ourselves and to others, we can avoid such bickering both within our own discipline and with those who view linguistics from a very different perspective.

Primary commitments

For me, cognitive linguistics is defined by two primary commitments, what I will call the Generalization Commitment and the Cognitive Commitment. The generalization commitment is a commitment to characterizing the general principles governing all aspects of human language. I see this as the commitment to undertake linguistics as a scientific endeavour. The cognitive commitment is a commitment to make one's account of human language accord with what is generally known about the mind and the brain, from other disciplines as well as our own.

The generalization commitment comes with a phenomenological characterization of subfields in terms of the kinds of generalizations required:

In syntax: Generalizations about the distribution of grammatical morphemes, categories, and constructions.
In semantics: Generalizations about inferences, polysemy, semantic fields, various kinds of semantic relationships, conceptual structure, knowledge structure, and the fitting of language to what we perceive, experience, and understand.
In pragmatics: Generalizations about speech acts, discourse, implicatures, deixis, and the use of language in context.

And so on, for morphology, phonology, etc. Of course, no a priori commitment is made as to whether these are separate subfields. It is an empirical matter, and empirical considerations suggest that they are not — that, for example, generalizations about syntax depend on semantic and pragmatic considerations.

The cognitive commitment forces one to be responsive to a wide variety of empirical results from a number of disciplines. Examples include:

Categorization results from cognitive psychology, developmental psy-

chology, and anthropology that demonstrate the existence of basic-level categorization and prototype effects.
Psychophysical, neurophysiological, anthropological results about the nature of color perception and categorization.
Results from cognitive psychology concerning human imaging capacities and the association of conventional imagery with language.
Results from cognitive neuroscience and connectionism regarding the computational mechanisms of the brain.

If we are fortunate, these commitments will mesh: the general principles we seek will be cognitively real. If not, the cognitive commitment takes priority: we are concerned with cognitively real generalizations.

This is anything but a trivial matter. Cognitively real generalizations may not at all accord with generalizations arrived at by classical techniques of linguistic analysis. For example, the fact that cognitive categories are, for the most part, not classical means that a linguistic analyst who defines what a generalization is only in terms of classical categories (given, say, by lists of features) will inevitably be in conflict with the cognitive commitment. The cognitive theory of categorization is presupposed in characterizing what counts as a generalization in cognitive linguistics. To a generative linguist, classical categories are the only ones possible and generalizations must be defined using them. A cognitive linguist, on the other hand, expects categories to have one of the various kinds of prototype structures (Lakoff 1987) and to be organized in terms of basic-level, superordinate and subordinate levels. Within cognitive linguistics, the use of classical categories in an analysis cries out for empirical justification: a demonstration that there is no prototype or basic-level structure.

The cognitive commitment also leads one to be suspicious of unstructured lists. One thing that we know is that the warehouse theory of memory is wrong: neural networks do not just learn isolated pieces of information about a subject matter without generalizing (Rumelhart and McClelland 1986: chap. 14). Thus we do not expect a linguistic system to be an unstructured list of lexical items and constructions: we expect to find networks of relationships among lexical items and constructions.

Taking the generalization and cognitive commitments as primary makes all other commitments secondary: philosophical commitments, commitments as to the proper form of linguistic descriptions, and other assumptions about the nature of reason and language. The import of this is that, when primary and secondary commitments come into conflict with one another, the primary commitment wins out and the secondary commitment must be given up.

Let me take an example. Back in the days when I was a generative linguist, my commitment to the symbol-manipulation paradigm that defines generative linguistics was secondary. I maintained that commitment for many years, so long as it did not contradict the primary generalization and cognitive commitments that have always defined my work. But the discovery in the mid-1970s of basic-level and prototype-based categorization and the subsequent establishment of the need for image–schemas in order to characterize certain linguistic generalizations created a crisis of commitment for me. The cognitive commitment required that I take basic-level categorization seriously. The generalization commitment required that I take image-schemas seriously. Since they could not be accounted for in the generative symbol-manipulation paradigm, and since the generative commitment was secondary, I gave up the generative commitment. Though the generative commitment is *logically* consistent with the cognitive and generalization commitments, it is *empirically* inconsistent with them. To maintain the generative commitment would have meant giving up on the cognitive and generalization commitments, that is, giving up on cognitive linguistics as I understand it.

Similarly, I used to have another secondary commitment, what I will call the Fregean commitment to a view of meaning based on truth and reference. This led, in the early 1960s, to my suggesting that such mechanisms of logic as logical form and model theory were needed in linguistics. However, in the 1970s, it became clear to me that the Fregean commitment was empirically inconsistent with the cognitive and generalization commitments (Lakoff 1989: 55–76). What made that clear was not only the discovery of basic-level categories and image-schemas, but also the discovery of conceptual metaphor. The theory of conceptual metaphor is an empirical consequence of applying the generalization commitment to the phenomena of polysemy and inference (see Lakoff and Brugman, 1986). Without conceptual metaphor, a large range of generalizations cannot be stated. Maintaining a commitment to stating such generalizations means giving up on the Fregean commitment, that is, giving up on the apparatus of formal logic and on the idea that meaning is based on reference and truth.

The cognitive commitment requires us to take research in cognitive psychology seriously, and hence also motivates a theory of conceptual metaphor. Recent research by Ray Gibbs and his co-workers at Santa Cruz has experimentally confirmed our conclusions about the existence of conceptual metaphor and conventional mental images.

In short, accepting the generalization and cognitive commitments is no small matter. These commitments have far-reaching consequences when combined with empirical research — consequences that utterly change the

nature of linguistics. Thus, my present views on metaphor, image-schemas, radial categories and prototype theory in general are not a priori commitments in themselves. They are *empirical consequences* of adopting the generalization and cognitive commitments — what I take as defining the field of cognitive linguistics.

I have endeavored to make my own commitments clear for a number of reasons:

First, everything I say will be based on those commitments. Since some of what I say is bound to be controversial, I want to factor out controversy about commitments from controversy about analyses (where a commonality of commitments is assumed).

Second, I do not know how many scholars who consider themselves cognitive linguists also hold the same commitments I do, but the only way to find out is to state our commitments overtly.

Third, I am sure that others who consider themselves cognitive linguists do not have the same primary commitments that I do, and that disagreements over how to properly analyze a given phenomenon are sure to follow from differences in primary commitments.

Cognitive vs. generative linguistics

As I understand it, generative linguistics is defined by the primacy of the generative commitment: the commitment to view language in terms of systems of combinatorial mathematics of the sort first characterized by the mathematician Emil Post. Such systems are called "formal grammars". They are systems in which arbitrary symbols are manipulated by rules of a restricted mathematical form without taking into account the interpretation of those symbols.

What is excluded by the generative commitment is everything not characterizable in terms of such systems: mental images and image-schemas, general cognitive processes, basic-level categories (which are defined partly in sensorimotor terms), prototype phenomena in general, the meanings of the symbols used, the grounding of meaning in bodily and social experience, and the use of neural foundations for linguistic theory. Let us refer to these as "nonfinitary phenomena".

To accept the generative commitment as primary is to define the study of linguistics in terms of the study of formal grammars and hence to restrict the study of language to what such systems can do and to exclude from the study of language all of those things just mentioned that are excluded by the commitment to such systems. The autonomy of linguistics for generative linguists is thus not a consequence of anything empirical; it is rather a consequence of defining the field of linguistics in terms of the

primacy of the generative commitment to using formal grammars in the technical sense of the term. The mathematical properties of those grammars require autonomy since they cannot deal naturally with nonfinitary phenomena.

Taking the generative commitment as primary makes the generalization and cognitive commitments secondary. If linguistics is defined as exluding all nonfinitary phenomena, then any linguistic generalization that makes reference to nonfinitary phenomena will not count as a true linguistic generalization and so will be ignored (or most likely not even be noticed). Within generative linguistics, the very concept of a generalization is defined to a large extent by taking the generative commitment as primary. As to the cognitive commitment, it requires us not to ignore nonfinitary phenomena; indeed, it requires us to pay special attention to them.

As I mentioned, it is logically possible for cognitive and generative linguistics to be the same enterprise: If linguistic generalizations never made reference to any nonfinitary phenomena, if every aspect of language really were fully and adequately characterizable in terms of combinatorial systems, then cognitive and generative linguistics would be identical. But it is an empirical fact that they are not identical, that general principles of language not only make use of nonfinitary phenomena, but do so in virtually every aspect of their structure. It is empirical observation that has given rise to cognitive linguistics — the hundreds, perhaps even thousands, of cases described so far where those phenomena excluded by generative linguistics are needed to state the general principles governing language.

It is not merely the case that cognitive linguistics covers more phenomena than generative linguistics. It does that, but it covers those phenomena in a very different way. Take, for example, the nature of semantic representation. The cognitive and generalization commitments have led cognitive linguists to hypothesize notions like image–schemas, metaphoric and metonymic mappings, mental spaces, radial categories and so on in order to characterize semantic generalizations. The phenomena that have led to such conclusions are usually not discussed by generative linguists, primarily, I think, because the descriptive apparatus available to generative linguists is not capable of stating general principles governing such phenomena. This is of course not seen as problematic for generative linguists because their discipline is defined in a restricted way so as to exclude those phenomena.

In addition to this empirically-based distinction between the disciplines, there is also a philosophical distinction, a distinction about what counts as knowledge and as sound scientific practice. Both enterprises see themselves as scientific and as committed to maximizing precision. But generative linguists tend to define precision as the use of the mathematics of combinatorial systems, while cognitive linguists have no such restric-

tions on what counts as precision. Thus, when Noam Chomsky described generative linguistics as committed to no more than being "precise and complete", he was assuming that the use of certain systems of combinatorial mathematics was the only way to be "precise". What was to count as "complete" was thereby relativized to what was to count as "precise": General principles not describable in terms of formal grammars were not taken as true linguistic principles and hence were not required to be part of a "complete" description of a language. Given this commitment as to what counts as precise and, hence, scientific, only generative linguistics is seen as "scientific".

Cognitive linguistics has a very different view as to what counts as scientific: To those who take the generalization and cognitive commitments as primary, the scientific study of language consists in seeking general principles governing all of language consistent with our overall knowledge about cognition and the brain. From a cognitive perspective, taking the generative commitment as primary appears unscientific because it excludes, a priori, the study of all linguistic regularities that cannot be expressed in the form of certain combinatorial mathematical systems. This is an empirically arbitrary restriction that makes it impossible to state overall general principles governing all aspects of language. From this perspective, generative linguistics appears as a philosophical program rather than a scientific enterprise — the study of the consequences of taking the generative commitment as primary.

One can now see why cognitive and generative linguists often have problems communicating with each other. They take different commitments as primary. These initial commitments are not only empirically incompatible on a massive scale, but they also entail very different views of what linguistics is as a scientific enterprise. Given such differences, it would be miraculous if communication were easy. Communication will be possible only if such differences in initial commitments are recognized.

It should be noted that not everyone who identifies himself professionally as a "generative linguist" undertakes what I have called the generative commitment. Many linguists take "precise" and "complete" to be English words rather than technical terms and accept a broadly and non-technically construed commitment to precision and completeness. Such a commitment is, of course, consistent with cognitive linguistics. But it is far from the technical commitment made by mainstream generative theoreticians.

Varieties of cognitive linguistics

Those who identify themselves professionally as cognitive linguists will not necessarily share the initial commitments that I see as defining the

cognitive linguistics enterprise. Let me take an obvious example: Professor Anna Wierzbicka has long maintained a primary commitment to the existence of a universal set of semantic primitives of the sort suggested by Leibniz. Her commitment to such a specific view of "The Alphabet of Human Thoughts" takes priority for her over what I have called the generalization and cognitive commitments. Because we start out with different primary commitments, Professor Wierzbicka and I are likely to disagree on many matters. Her primary commitment is inconsistent with the theory of conceptual metaphor, prototype theory, the theory of basic-level concepts, etc. It is therefore inevitable that we will disagree on these and other matters.

I assume that Professor Wierzbicka would not accept the generalization and cognitive commitments as taking precedence over her Leibnizian commitment. I would therefore not expect her to give up her Leibnizian commitment in the face of counterevidence of the sort I have presented in various works. Given such initial differences, it is inevitable that we will reach different conclusions. Without agreement on initial premises, arguments about conclusions will be pointless.

I have singled out Professor Wierzbicka because of my great respect for her distinguished contributions to linguistics over a long career and because I have learned much from her work despite our disagreements. Similar differences in primary commitments are inevitable among members of the cognitive linguistic community, as they are in any scientific community. What is important is that we understand the nature of such disagreements, that we acknowledge them overtly, and that we be able to discuss them openly without rancor.

There is, of course, a very good reason why I have taken the generalization and cognitive commitments as primary. From my perspective, the generalization commitment is a commitment to linguistics as a scientific endeavor, a commitment to seek general principles. The cognitive commitment is a commitment not to isolate linguistics from the study of the mind, but to take seriously the widest range of other data about the mind. Neither of these commitments, in themselves, imposes a particular form on the answer. As such, they are methodological, not substantive commitments.

By contrast, the generative, Fregean, and even the Leibnizian commitments all presuppose the form of an answer. The generative commitment requires an answer in terms of the manipulation of uninterpreted symbols. The Fregean commitment requires an answer in the form of truth conditions and mappings from symbols to things in the world. The Leibnizian commitment requires an answer in the form of a small alphabet of primitives.

I prefer the primary commitments that I and many other cognitive linguists have made because they do not impose a particular form on the answer and so do not artificially limit the inquiry. What has been interesting about them is that they have led, on empirical grounds, to a very rich conception of the nature of language and thought.

2. Some basic properties of metaphor

The theory of metaphor, as it has evolved in the past decade, arises from the application of the generalization and cognitive commitments to a wide range of data. To take an example, English is full of expressions that reflect the conceptualization of love as a journey. Some are necessarily about love; others can be understood that way:

Look *how far we've come.*
It's been *a long, bumpy road.*
We can't *turn back* now.
We're at a *crossroads.*
We may have to *go our separate ways.*
We're *spinning our wheels.*
The relationship isn't *going anywhere.*
The marriage is *on the rocks.*

These are ordinary, everyday English expressions. They are not poetic, nor are they necessarily used for special rhetorical effect. Those like *Look how far we've come*, which are not necessarily about love, can readily be understood as being about love. Examples like this show that what is involved is not just conventional language, but a conventional mode of thought. They reflect a way of thinking about love in terms of a certain kind of journey:

The lovers are travellers on a journey together, with their common life goals seen as destinations to be reached. The relationship is their vehicle, and it allows them to pursue those common goals together. The relationship is seen as fulfilling its purpose as long as it allows them to make progress toward their common goals. The journey is not easy. There are impediments, and there are places (crossroads) where a decision has to be made about which direction to go and whether to keep travelling together. The mode of travel can be of various types: car (*long bumpy road, spinning our wheels*), train (*off the track*), boat (*on the rocks, foundering*), plane (*just taking off, bailing out*).

The metaphor involves understanding one domain of experience, love, in terms of a very different domain of experience, journeys. The metaphor

can be understood as a mapping (in the mathematical sense) from a source domain (in this case, journeys) to a target domain (in this case, love). The mapping is tightly structured. There are ontological correspondences, according to which entities in the domain of love (for example, the lovers, their common goals, their difficulties, the love relationship, etc.) correspond systematically to entities in the domain of a journey (the travellers, the vehicle, destinations, etc.). Some examples of ontological correspondences are the following:

The lovers correspond to travellers.
The love relationship corresponds to the vehicle.
The state of being in the relationship corresponds to travelling in the same vehicle.
The intimacy of being in the relationship corresponds to the physical closeness of being in the vehicle.
The lovers' common goals correspond to their common destinations on the journey.
Difficulties correspond to impediments to travel.

The mapping includes epistemic correspondences, in which knowledge about journeys is mapped onto knowledge about love. Such correspondences permit us to reason about love using the knowledge we use to reason about journeys. Let us take an example:

Two travellers are travelling somewhere in a vehicle and it hits some impediment and gets stuck. If they do nothing, they will not reach their destinations.

There are a limited number of alternatives for action.

1. They can try to get it moving again, either by fixing it or getting it past the impediment that stopped it.
2. They can remain in the stuck vehicle, and give up on getting to their destinations in it.
3. They can abandon the vehicle.

The alternative of remaining in the stuck vehicle takes the least effort, but does not satisfy the desire to reach their destinations.

The ontological correspondences map this scenario (sometimes called a "knowledge structure" in the cognitive sciences) onto a corresponding love scenario, in which the corresponding alternatives for action are seen. Here is the corresponding love scenario that results from applying the ontological correspondences to this knowledge structure.

Two people are in love and pursuing their common goals in a love relationship. They encounter some difficulty in the relationship which, if

nothing is done, will keep them from pursuing their goals. Here are their alternatives for action:

1. They can try to do something so that the relationship will once more allow them to pursue their goals.
2. They can leave the relationship as it is and give up on pursuing those goals.
3. They can abandon the relationship.

The alternative of remaining in the relationship takes the least effort, but does not satisfy goals external to the relationship.

What constitutes the love-as-journey metaphor is not any particular word or expression. It is the ontological and epistemic mapping across conceptual domains, from the source domain of journeys to the target domain of love. The metaphor is not just a matter of language, but of thought and reason. The language is a reflection of the mapping. The mapping is conventional, one of our conventional ways of understanding love.

If metaphors were just linguistic expressions, we would expect different linguistic expressions to be different metaphors. Thus, "We've hit a dead-end street" would constitute one metaphor. "We can't turn back now" would constitute another, quite different metaphor. 'Their marriage is on the rocks' would involve a still different metaphor. And so on for dozens of examples. Yet we do not seem to have dozens of different metaphors here. We have one metaphor, in which love is conceptualized as a journey. It is a unified way of *conceptualizing* love metaphorically that is realized in many different *linguistic* expressions.

It is here that the generalization commitment comes into play. The love-as-journey metaphor characterizes a linguistic generalization of two kinds:

Polysemy generalization: A generalization over related senses of linguistic expressions, for example, *dead-end street*, *crossroads*, *spinning one's wheels*, *not going anywhere*, and so on.
Inferential generalization: A generalization over inferences across different conceptual domains.

As long as the generalization commitment is among one's primary commitments, then such evidence will lead to the conclusion that there is conceptual metaphor.

There is, of course, further evidence that can be brought to bear for such a conceptual metaphor, evidence arising from the predicative value of the metaphor. Take a sentence like:

Look how far we've come.

This can be about love (as well as about other activities, say careers, that

are conceptualized as journeys). None of the individual words would be listed in an English lexicon as being about love. Not "look" or "far" or "come" (in the sense used here). Such a fact can, however, be explained if we hypothesize a conceptual love-as-journey metaphor.

Such a conceptual metaphor explains why new and imaginative extensions of the mapping can be understood instantly, given the ontological correspondences and other knowledge about journeys. Take the song lyric,

We're driving in the fast lane on the freeway of love.

The travelling knowledge called upon is this: When you drive in the fast lane, you go a long way in a short time and it can be exciting and dangerous. The general metaphorical mapping maps this knowledge about driving into knowledge about love relationships. The danger may be to the vehicle (the relationship may not last) or the passengers (the lovers may be hurt, emotionally). The excitement of the love-journey is sexual. Our understanding of the song lyric depends upon the pre-existing metaphorical correspondences of the love-as-journey metaphor. The song lyric is instantly comprehensible to speakers of English because those metaphorical correspondences are already part of our conceptual system.

The love-as-journey metaphor was the example that first convinced me that metaphor was not a figure of speech, but a mode of thought, defined by a systematic mapping from a source to a target domain. What convinced me were the three characteristics of metaphor that we have just discussed:

1. The systematicity in the linguistic correspondences.
2. The use of metaphor to govern reasoning and behavior based on that reasoning.
3. The possibility for understanding novel extensions in terms of the conventional correspondences.

So far, we have discussed only the generalization commitment. Let us turn to the cognitive commitment, which commits one to taking experimental results from cognitive psychology seriously. In the past, I have held to this commitment in adjusting my theoretical views to fit experimental results on the nature of categorization, both prototype and basic-level results. But what about metaphor?

Many of the metaphorical expressions discussed in the metaphor literature are idioms. On classical views, idioms have arbitrary meanings. But within cognitive linguistics, the possibility exists that they are not arbitrary, but rather motivated, and conceptual metaphor can be one of the things motivating an idiom. Let us look a little more closely at idioms.

An idiom like "spinning one's wheels" comes with a conventional mental image, that of the wheels of a car stuck in some substance — either in mud, sand, snow, or on ice, so that the car cannot move when the motor is engaged and the wheels turn. Part of our knowledge about that image is that a lot of energy is being used up (in spinning the wheels) without any progress being made, that the situation will not readily change of its own accord, that it will take a lot of effort on the part of the occupants to get the vehicle moving again — and that may not even be possible.

The love-as-journey metaphor applies to this knowledge about the image associated with "spinning one's wheels" to map this knowledge about cars onto knowledge about love relationships: A lot of energy is being spent without any progress toward fulfilling common goals, the situation will not change of its own accord, it will take a lot of effort on the part of the lovers to make more progress, and so on. In short, when idioms have associated conventional images, it is common for an independently-motivated conceptual metaphor to map that knowledge from the source to the target domain. At least, this is what one is led to by the generalization commitment. Cognitive psychologists Ray Gibbs and Jennifer O'Brien at the University of California at Santa Cruz have run three classes of experiments to test this analysis. Their experiments confirm such analyses overwhelmingly (see Gibbs and O'Brien 1989). In such cases, the cognitive commitment and the generalization commitment lead one to the same conclusions.

The moral: If one accepts the cognitive and generalization commitment as primary, then one must accept the account of metaphor and of imageable idioms that they entail. These conclusions can be avoided only by placing some other commitments ahead of the cognitive and generalization commitments.

3. The metaphorical understanding of basic semantic concepts

Most people are not too surprised to discover that emotional concepts like love and anger are understood metaphorically. What is more interesting, and I think more exciting, is the realization that many of the most basic concepts in semantics are also understood metaphorically — concepts like time, quantity, state, change, action, cause, purpose, means, modality and even the concept of a category. These are concepts that enter normally into the grammars of languages, and if they are indeed metaphorical in nature, then metaphor becomes central to grammar.

What I would like to suggest is that the same kinds of considerations that lead to our acceptance of the love-as-journey metaphor lead inevita-

bly to the conclusion that such basic concepts are often, and perhaps always, understood via metaphor.

Categories

Classical categories are understood metaphorically in terms of bounded regions, or "containers". Thus, something can be *in* or *out* of a category, it can be *put into* a category or *removed from* a category, etc. The logic of classical categories is the logic of containers (see Figure 1).

If X is in container A and container A is in container B, then X is in container B.

This is true not by virtue of any logical deduction, but by virtue of the topological properties of containers. Under the CLASSICAL CATEGORIES ARE CONTAINERS metaphor, the logical properties of categories are inherited from the logical properties of containers. One of the principal logical properties of classical categories is that the classical syllogism holds for them. The classical syllogism,

Socrates is a man.
All men are mortal.
Therefore, Socrates is mortal.
is of the form:

If X is in category A, and category A is in category B, then X is in category B.

Thus, the logical properties of classical categories can be seen as following from the topological properties of containers plus the metaphorical mapping from containers to categories. As long as the topological

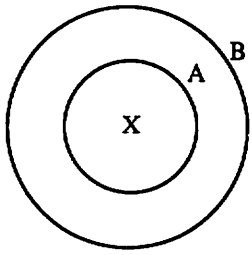

X is in A
A is in B
∴ X is in B

Figure 1.

properties of containers are preserved by the mapping, this result will be true.

In other words, there is a generalization to be stated here. The language of containers applies to classical categories and the logic of containers is true of classical categories. A single metaphorical mapping ought to characterize both the linguistic and logical generalizations at once. This can be done provided that the topological properties of containers are preserved in the mapping.

The joint linguistic-and-inferential relation between containers and classical categories is not an isolated case. Let us take another example.

Quantity and linear scales

The concept of quantities involves at least two metaphors. The first is the well-known MORE IS UP; LESS IS DOWN, as shown by a myriad of expressions like *Prices rose, Stocks skyrocketed, The market plummeted,* and so on. A second is that LINEAR SCALES ARE PATHS. We can see this in expressions like:

John is *far* more intelligent than Bill.
John's intelligence *goes way beyond* Bill's.
John is *way ahead of Bill in intelligence.*

The metaphor maps the starting point of the path onto the bottom of the scale and maps distance traveled onto quantity in general.

What is particularly interesting is that the logic of paths maps onto the logic of linear scales. (See Figure 2.)

Path inference: If you are going from A to C, and you are now at an intermediate point B, then you have been at all points between A and B and not at any points between B and C.

Example: If you are going from San Francisco to N.Y. along route 80, and you are now at Chicago, then you have been to Denver but not to Pittsburgh.

Linear scale inference: If you have exactly $50 in your bank account, then you have $40, $30, and so on, but not $60, $70, or any larger amount.

The form of these inferences is the same. The path inference is a consequence of the cognitive topology of paths. It will be true of any path image-schema. Again, there is a linguistic-and-inferential generalization to be stated. It would be stated by the metaphor LINEAR SCALES ARE PATHS, provided that metaphors in general preserve the cognitive topology (that is, the image-schematic structure) of the source domain.

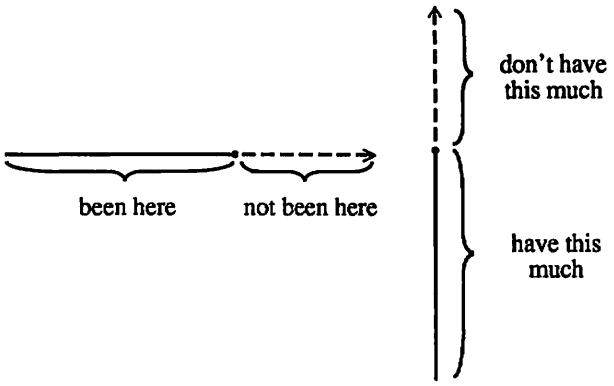

Figure 2.

The Invariance Hypothesis

Such considerations have led to the following hypothesis:

The Invariance Hypothesis: Metaphorical mappings preserve the cognitive topology (this is, the image–schema structure) of the source domain.

It will follow from this that all source domain inferences due to cognitive topology (image–schema structure) will be preserved in the mapping. This would account for what has been observed empirically in metaphor studies up to the present, that metaphors preserve inferential structure — at least certain kinds of inferential structure.

It would also follow from this that a great many, if not all, abstract inferences are actually metaphorical versions of spatial inferences that are inherent in the topological structure of image–schemas. Thus, the generalization commitment leads us to the Invariance Hypothesis, which in turn makes a major controversial claim about the nature of abstract reason.

The Invariance Hypothesis has another consequence as well, a consequence for the type of imagistic representation that Ron Langacker has proposed for many abstract concepts. The Invariance Hypothesis claims that, if those abstract concepts are metaphorically understood, then their imagistic representations are the image–schemas that have been metaphorically projected from the source domains of the metaphors. In short, the Invariance Hypothesis is a possible link between metaphor and Langacker–style analysis. Indeed, that link appears to hold for the examples we have discussed so far, classical categories and linear scales.

What I will do now is turn to other cases of basic, but abstract,

semantic concepts to see what evidence there is for a metaphoric understanding of such concepts, and then return afterward to the question of the Invariance Hypothesis.

Time

It has often been noted that time in English is conceptualized in terms of space. The details are rather interesting.

Ontology: Time is understood in terms of things (i.e., entities and locations) and motion.
Background condition: The present time is at the same location as a canonical observer.

Mapping:
Times are things.
The passing of time is motion.
Future times are in front of the observer; past times are behind the observer.
One thing is moving, the other is stationary; the stationary entity is the deictic center.

Entailment:
Since motion is continuous and one-dimensional, the passage of time is continuous and one-dimensional.

Special case 1:
The observer is fixed; times are entities moving with respect to the observer.
Times are oriented with their fronts in their direction of motion.

Entailments:
If time 1 follows time 2, then time 1 is in the future relative to time 2.
The time passing the observer is the present time.
Time has a velocity relative to the observer.

Special case 2:
Times are fixed locations; the observer is moving with respect to time.

Entailment:
Time has extension, and can be measured.
An extended time, like a spatial area, may be conceived of as a bounded region.

This metaphor, with its two special cases, embodies a generalization that accounts for a wide range of cases where a spatial expression can also be used for time. Special case 1 accounts for both the linguistic form and the semantic entailments of expressions like:

The time will come when ...
The time has long since gone when ...
The time for action has arrived.
That time is here.
In the weeks following next Tuesday ...
On the preceding day, ...
I'm looking ahead to Christmas.
Thanksgiving is coming up on us.
Let's put all that behind us.
I can't face the future.
Time is flying by.
The time has passed when ...

Thus, special case 1 characterizes the general principle behind the temporal use of words like *come*, *go*, *here*, *follow*, *precede*, *ahead*, *behind*, *fly*, *pass*, accounting not only for why they are used for both space and time, but why they mean what they mean.

Special case 2 accounts for a different range of cases, expressions like:

There's going to be trouble down the road.
He stayed there for ten years.
He stayed there a long time.
His stay in Russia extended over many years.
He passed the time happily.
He arrived on time.
We're coming up on Christmas.
We're getting close to Christmas.
He'll have his degree within two years.
I'll be there in a minute.

Special case 2 maps location expressions like *down the road*, *for* + location, *long*, *over*, *come*, *close to*, *within*, *in*, *pass*, onto corresponding temporal expressions with their corresponding meanings. Again, special case 2 states a general principle relating spatial terms and inference patterns to temporal terms and inference patterns.

The details of the two special cases are rather different; indeed, they are inconsistent with one another, which is what makes them special cases. The existence of such special cases has an especially interesting theoretical consequence: words mapped by both special cases will have inconsistent readings. Take, for example, the *come* of *Christmas is coming* (special case 1) and *We're coming up on Christmas* (special case 2). Both instances of *come* are temporal, but one takes a moving time as first argument and the other takes a moving observer as first argument. The same is true of *pass*

in *The time has passed* (special case 1) and in *He passed the time* (special case 2).

These differences in the details of the mappings show that one cannot just say blithely that spatial expressions can be used to speak of time, without specifying details, as though there were only one correspondence between time and space. When we are explicit about stating the mappings, we discover that there are two different — and inconsistent — subcases.

The fact that time is understood metaphorically in terms of motion, entities, and locations accords with our biological knowledge. In our visual systems, we have detectors for motion and detectors for objects/locations. We do not have detectors for time (whatever that could mean). Thus, it makes good biological sense that time should be understood in terms of things and motion.

Does it follow that time is never understood in its own terms, with some structure independent of metaphor? The answer is no. We have no evidence one way or the other. There could be some structure in the time domain that is independent of any metaphor and neutral between the two special cases of basic time metaphor, that is, a structure sufficiently underspecified that the two special cases can both map onto it. But in any given sentence, one of the two special cases of the time metaphor may be imposing its structure. Thus, we cannot simply give a single undifferentiated temporal analysis for a given sentence. Rather, we need to keep track of the special cases of the time metaphor to see if either is present in a given sentence.

Event structure

This is a report on some as yet unpublished work by myself and two of my students, Sharon Fischler and Karin Myhre, on what we have discovered about the metaphorical understanding of event structure in English. What we have found is that various aspects of event structure, including notions like states, changes, processes, actions, causes, purposes, and means, are understood metaphorically in terms of space, motion, and force.

The general mapping we have found goes as follows:

States are bounded regions in space.
Changes are movements into or out of bounded regions.
Processes are movements.
Actions are self-propelled movements.
Causes are forces.
Purposes are destinations.
Means are paths to destinations.

This mapping generalizes over an extremely wide range of expressions for one or more aspects of event structure. For example, take states and changes. We speak of being *in* or *out* of a state, of *going into* or *out of* it, of *entering* or *leaving* it, of getting *to* a state or emerging *from* it.

To see just how rich this metaphor is, consider some of its basic correspondences:

Impediments to action are impediments to motion.
Manner of action is manner of motion.
A different means for achieving a purpose is a different path.
Forces affecting action are forces affecting motion.
The inability to act is the inability to move.
Progress made is distance travelled.

We will consider examples of each of these one-by-one, including a number of special cases.

Impediments to Action are Impediments to Motion
We hit a roadblock.
We are at an impasse.
I can't find my way around that.
I've hit a brick wall.
We are going upstream.
We are fighting an uphill battle.
It's a steep road ahead.
It's a long and winding road.
We are in rough waters.

Aids to Action are Aids to Motion
It is smooth sailing from here on in.
It's all downhill from here.
There's nothing in our way.

A Different Means of Achieving a Result is a Different Path
Do it this way.
She did it the other way.
Do it any way you can.
However you want to go about it is fine with me.

Forced Motion is Forced Action
He pushes me too hard.
She pushed me into doing it.
They dragged me into doing it.
I am being pulled along by the current.
She leaned on him to do it.

She put the lean on him.
He is a mover and a shaker.
He really throws his weight around.

Guided Action is Guided Motion
She guided him through it.
She walked him through it.
She led him through the rough parts.

Inability to Act is Inability to Move
We are stuck on this problem.
I am drowning in work.
I am tied up with work.
He is up to his neck in work.

Special case: Suspension (If one is hanging above the path, then one cannot move along it)
I am really hung up on this problem.
He is so caught up in his work he can't do anything else.
He was held up in the meeting.
He was hung up at school.

A Force That Limits Action Is A Force That Limits Motion
She leads him around by the nose.
She held him back.
She is being pushed into a corner.
He is up against a wall.
I am being pulled under.
He doesn't give me any slack.
She has him on a tight rein.
She has him on a short leash.
He is tied to his mother's apron strings.
He is tied up with work.

Manner of Action is Manner of Motion
We are moving/running/skipping right along.
We slogged through it.
He is flailing around.
He is falling all over himself.
We are leaping over hurdles.
He is out of step.
He is in step.

Careful Action is Careful Motion
I'm walking on eggshells.
He is treading on thin ice.
He is walking a fine line.

Speed of Action is speed of Movement
He flew through his work.
He is running around.
I have been running all day.
It is going swimmingly.
Keep things moving at a good clip.
Things have slowed to a crawl.
She is going by leaps and bounds.
I am stagnating.
I am moving at a snail's pace.

Purposeful Action is Directed Motion To a Destination
This has the following special cases:

Progress Is Forward Movement
We are moving ahead.
Let's forge ahead.
Let's keep moving forward.
We made lots of forward movement.

Progress is Distance Moved
We've come a long ways.
We've covered lots of ground.
We've made it this far!

Undoing Progress is Backward Movement
We are sliding backward.
We are backsliding.
We need to backtrack.
It is time to turn around and retrace our steps.

Starting an Action is Starting out on a Path
We are just starting out.
We have taken the first step.

Success Is Reaching The End of the Path
We've reached the end.
We are seeing the light at the end of the tunnel.
We only have a short way to go.
The end is in sight.
The end is a long ways off.

Lack of Purpose is Lack of Direction
He is just floating around.
He is drifting aimlessly.
He needs some direction.

Lack of Progress is Lack of Movement
We are at a standstill.
We aren't getting any place.
We aren't going anywhere.
We are going nowhere with this.

These examples show that the event structure metaphor exists and that it does a lot of work in characterizing how all of these expressions involving space, motion, and force can be used to talk and reason about states, events, actions, causes, purposes, and means. However, there is more metaphorical complexity to causation and to change than we have seen thus far.

Let us begin with change. Ken Baldwin, in an unpublished study of the verb *turn*, has observed that in expressions like

The milk turned sour.

there is another metaphor for change, one that concerns the maintenance or change of state over time:

Change of state is change of direction.
Maintenance of state is maintenance of direction.

Examples of the second part of the metaphor are:

We're in a rut.
Things are going the way they've always gone.

Putting this together with the causes-as-forces portion of the event structure metaphor, we get the entailment that:

Causing a change of state is forcing a change of direction.
This is the metaphor behind such expressions as:

We need to take the country in a new direction.
We're going to move the country down the path to a drug-free society.

Causation

In the example just given, the causes-as-forces metaphor interacts with another metaphor for change of state to yield the complex result that causing a change of state is forcing a change of direction. Current research by Jane Espenson at Berkeley (personal communication) suggests that this kind of complex interaction is common for the causes-as-forces metaphor. Here is a brief summary of her results.

Caused action, as we have seen, is understood as forced motion. There

are two principal kinds of forced motion: propulsion (sending, throwing, propelling, etc.) and the continuous application of force to produce motion (as in bringing or giving). These have different entailments. With continuous application, motion continues only as long as the force is applied. With propulsion, the application of force begins the motion, which continues afterwards. These entailments about force are mapped onto causation via the CAUSES ARE FORCES metaphor.

Consider the following examples:

The home run brought the crowd to its feet.
The home run sent the crowd into a frenzy.

Here *bring* and *send* are both being used as causative verbs, since both, in their central senses, involve forced motion. But, because they involve different kinds of forced motion — continuous application versus propulsion — the CAUSES ARE FORCES metaphor maps them into different kinds of causation. In the first example with *brought*, the effect of the cause goes on during the flight of the ball and then ceases: the crowd rises to its feet while the ball is in the air. In the second case with *send*, the frenzy ensues after the home run. Thus, two special cases of force are mapped into two special cases of causation by the CAUSES ARE FORCES metaphor.

Espenson has further noted that CAUSES ARE FORCES interacts with other existing metaphors to yield an even richer variety of causation types. There is a metaphor in English to the effect that EXISTENCE IS LOCATION HERE; NONEXISTENCE IS LOCATION AWAY. Since change is motion to a bounded area, and existence is metaphorized as a bounded area around where we are, something can *come into* existence or *go out* of existence, with the choice between *come* and *go* being determined by deictic center (*here*). The *about* of *come about* indicates an area in the vicinity of a deictic center, typically the speaker, which, via this metaphor, indicates the domain of existence. In this metaphor, an event happens when it comes into the domain of existence. Thus, *The revolution came about* means that the revolution occurred.

Putting CAUSES ARE FORCES together with EXISTENCE IS LOCATION HERE, we get expressions like *bring into existence* and *bring about*, where *bring* indicates the continuous application of force. For example, in

The stock market crash brought about political instability.

political instability is seen as coming into existence under the force of the stock market crash. In this type of causation, the metaphorical force is applied to the EVENT, moving it into existence. This is rather different from the case discussed above, namely,

The home run brought the crowd to its feet.

where the force is applied to the PATIENT (the crowd), moving it to a new state (being on its feet). Thus, we have two general patterns of causation, so far:

Case 1: CAUSES ARE FORCES plus STATES ARE BOUNDED AREAS.

The force applies to the patient, moving it to a new state (a bounded area).

Case 2: CAUSES ARE FORCES plus EXISTENCE IS LOCATION HERE.
The force applies to the event, moving it into existence (a bounded area around us).

A minimal pair illustrating these two metaphorical versions of causation would be:

Case 1: He brought the water to a boil.
Case 2: He brought about the boiling of the water.

Further interactions are possible between CAUSES ARE FORCES and other metaphors, for instance, PROPERTIES ARE POSSESSIONS and EXPERIENCES ARE POSSESSIONS. Suppose, for example, that patience is one of Harry's properties and that his patience is a result of the practice of Zen meditation. Because PROPERTIES ARE POSSESSIONS, we can speak of Harry as *having* patience and as having *acquired* patience. We can describe how he acquired patience using the verb *give*, as in:

The practice of Zen mediation gave Harry patience.

which attributes a causal role to the practice of Zen meditation. The verb *give* can be used to express causation because, in its central sense, *give* denotes a possession transfer: force is applied to an entity moving it to a recipient, who then possesses it. In this example, patience is the entity that moves into Harry's possession and the cause, Zen meditation, is seen as a causal force.

In general, *give* is used as a causal verb when CAUSES ARE FORCES combined with some possession metaphor. For example, it might combine with EXPERIENCES ARE POSSESSIONS, to yield an example like:

Problem 3 gave Harry trouble.

Here problem 3 is the cause of Harry's experiencing trouble. Via

EXPERIENCES ARE POSSESSIONS, we can think of Harry as *having trouble*, and via CAUSES ARE FORCES we can think of the cause as a force that moves the trouble-entity into Harry's possession. This, then, sanctions the use of the verb *give*. Note that this is still another type of causation.

Case 3: CAUSES ARE FORCES plus EXPERIENCES ARE POSSESSIONS.

The force applies to the possession, moving it to the possessor.

Thus, we have seen cases where the force can apply to an event, to a patient, and to a possession.

These are just some of the metaphorical models of causation discovered by Espenson. There are two morals here:
First, causes (as Talmy has observed in his work on force dynamics), are understood metaphorically as forces. Thus, causation is not a semantically primitive notion, independent of all metaphor.
Second, despite the existence of a single metaphor for causes, the interaction of that metaphor with other metaphors yields an extremely complex and disparate class of overall CAUSAL EVENTS. Thus, one cannot assume that all causal events have the same structure. They differ by something as elementary as what the causal force is applied to — events, patients, properties, etc.

Invariance again

The metaphors just given primarily map three kinds of image-schemas: containers, paths, and force-images. However, because of the sub-cases and interactions, the details are intricate, to say the least. However, the Invariance Hypothesis does make claims in each case as to what image-schemas get mapped onto target domains. I will not go through the details here, but so far as I can see, the claims made about inferential structure are reasonable ones.

For example, the logic of force dynamics does seem to map, via CAUSES ARE FORCES, onto the logic of causation. The following are inferences from the logic of forces inherent in force dynamics:

A stationary object will move only when force is applied to it; without force, it will not move.
The application of force requires contact; thus, the applier of the force must be in spatial contiguity with the thing it moves.
The application of force temporally precedes motion, since inertia must be overcome before motion can take place.

These are among the classic inferential conditions on causation: spatial contiguity, temporal precedence, and that A caused B only if B wouldn't have happened without A.

At this point, I would like to take up the question of what else the Invariance Hypothesis would buy us. I will consider two cases that arose while Mark Turner and I were writing *More Than Cool Reason* (1989). The first concerns image-metaphors and the second, generic-level metaphors. But before I move on to those topics, I should point out an important consequence of invariance.

Johnson and I argued in *Metaphors We Live By* (1980) that a complex propositional structure could be mapped by metaphor onto another domain. The main example we gave was argument-as-war. Kövecses and I, in our analysis of anger metaphors, also argued that metaphors could map complex propositional structures. The Invariance Hypothesis does not deny this, but it puts those claims in a very different light. Complex propositional structures involve semantic notions like time, states, changes, causes, purposes, quantity scales, and categories. If all of these abstract concepts are understood metaphorically, then the Invariance Hypothesis claims that what we had called propositional structure is really image–schematic structure! In other words:

So-called propositional inferences arise from the inherent topological structure of the image–schemas mapped by metaphor onto concepts like time, states, changes, actions, causes, purposes, means, quantity, and categories.

The reason that I have taken the trouble to discuss all those abstract concepts is to demonstrate this consequence of the Invariance Hypothesis; namely, that what have been seen in the past as propositional inferences are really image-based inferences. If the Invariance Hypothesis is correct, it has a startling consequence:

Abstract reasoning is a special case of image-based reasoning.

Image-based reasoning is fundamental and abstract reasoning is image-based reasoning under a metaphorical projection to an abstract domain.

To look for independent confirmation of the Invariance Hypothesis, let us turn to image-metaphors.

4. Image metaphors

There is a class of metaphors that function to map one conventional mental image onto another. These contrast with the metaphors we have

discussed so far, each of which maps one conceptual domain onto another, often with many concepts in the source domain mapped onto many corresponding concepts in the target domain. Image-metaphors, by contrast, are 'one-shot' metaphors: they map only one image onto one other image.

Consider, for example, this poem from the Indian tradition:

Now women-rivers
belted with silver fish
move unhurried as women in love
at dawn after a night with their lovers
(Merwin and Masson 1981: 71)

Here the image of the slow, sinuous walk of an Indian woman is mapped onto the image of the slow, sinuous, shimmering flow of a river. The shimmering of a school of fish is imagined as the shimmering of the belt.

Metaphoric image-mappings work in just the same way as all other metaphoric mappings: by mapping the structure of one domain onto the structure of another. But here, the domains are conventional mental images. Take, for example, this line from André Breton:

My wife ... whose waist is an hourglass.

This is a superimposition of the image of an hourglass onto the image of a woman's waist by virtue of their common shape. As before, the metaphor is conceptual; it is not in the words themselves, but in the mental images. Here, we have a mental image of an hourglass and of a woman, and we map the middle of the hourglass onto the waist of the woman. Note that the words do not tell us which part of the hourglass to map onto the waist, or even that it is only part of the hourglass shape that corresponds to the waist. The words are prompts for us to perform mapping from one conventional image to another. Similarly, consider:

His toes were like the keyboard of a spinet.
(Rabelais, 'The Descriptions of King Lent', trans. J.M. Cohen)

Here, too, the words do not tell us that an individual toe corresponds to an individual key on the keyboard. Again, the words are prompts for us to perform a conceptual mapping between conventional mental images. In particular, we map aspects of the part-whole structure of one image onto aspects of the part-whole structure of another. Just as individual keys are parts of the whole keyboard, so individual toes are parts of the whole foot.

Image-mapping can involve more than mapping physical part-whole relationships. For example, the water line of a river may drop slowly and

that slowness is part of the dynamic image, which may be mapped onto the slow removal of clothing:

Slowly slowly rivers in autumn show
sand banks
bashful in first love woman
showing thighs
(Merwin and Masson 1981: 69)

Other attributes are also mapped: the color of the sand bank onto the color of flesh, the quality of light on a wet sand bank onto the reflectiveness of skin, the light grazing of the water's touch receding down the bank onto the light grazing of the clothing along the skin. Notice that the words do not tell us that any clothing is involved. We get that from a conventional mental image. Part-whole structure is also mapped in this example. The water covers the hidden part of the bank just as the clothing covers the hidden part of the body. The proliferation of detail in the images limits image-mappings to highly specific cases. That is what makes them "one-shot" mappings.

Such mappings of one image onto another can lead us to map knowledge about the first image onto knowledge about the second. Consider the following example from the Navaho:

My horse with a mane made of short rainbows.
('War God's Horse Song I' Words by Tall Kia ahni.
Interpreted by Louis Watchman.)

The structure of a rainbow, its band of curved lines for example, is mapped onto an arc of curved hair, and many rainbows onto many such arcs on the horse's mane. Such image-mapping prompts us to map our evaluation of the source domain onto the target. We know that rainbows are beautiful, special, inspiring, larger than life, almost mystic, and that seeing them makes us happy and awe-inspired. This knowledge is mapped onto what we know of the horse: it too is awe-inspiring, beautiful, larger than life, almost mystic. This line comes from a poem containing a series of such image-mappings:

My horse with a hoof like a striped agate,
with his fetlock like a fine eagle plume:
my horse whose legs are like quick lightning
whose body is an eagle-plumed arrow:
my horse whose tail is like a trailing black cloud.

Image-metaphors raise two major issues for the general theory of metaphor:

How do they work? What constrains the mappings? What kind of internal structures do mental images have that permits some mappings to work readily, others only with effort, and others not at all?
What is the general theory of metaphor that unifies image-metaphors with all the conventional metaphors that map the propositional structure of one domain onto the propositional structure of another domain?

Turner and I (1989) have suggested that the Invariance Hypothesis could be an answer to both questions. We suggest that conventional mental images are structured by image–schemas and that image-metaphors preserve image–schematic structure, mapping parts onto parts and wholes onto wholes, containers onto containers, paths onto paths, and so on. The generalization would be that all metaphors are invariant with respect to their cognitive topology, that is, each metaphorical mapping preserves image–schema structure.

5. Generic-level metaphors

When Turner and I were writing *More Than Cool Reason* (1989), we hypothesized the existence of what we called "generic-level metaphors" to deal with two problems that we faced.

Problem 1: Personification

In studying a wide variety of poems about death in English, we found that, in poem after poem, death was personified in a relatively small number of ways: drivers, coachmen, footmen; reapers, devourers and destroyers; or opponents in a struggle or game (say, a knight or a chess opponent). The question we asked was: Why these? Why is death not personified as a teacher or a carpenter or an ice cream salesman? Somehow, the ones that occur repeatedly seem appropriate. Why?

In studying personifications in general, we found that the overwhelming number seem to fit a single pattern: events (like death) are understood in terms of actions by some agent (like reaping). It is that agent that is personified.

We thus hypothesized a very general metaphor, EVENTS ARE ACTIONS, in an attempt to make sense of these cases. But this metaphor was unlike any we had ever seen before: it was too general, it had no specific ontology, no specific mapping details. It also did not explain what could not be a personification of death.

What it did do was begin to make sense of some of the cases we had found when it was combined with other metaphors for life and death. Take, for example, the DEATH IS DEPARTURE metaphor. Departure is an event. If we understand this event as an action on the part of some causal agent — someone who brings about, or helps to bring about, departure — then we can account for figures like drivers, coachmen, footmen, etc. Or take the PEOPLE ARE PLANTS metaphor. In the natural course of things, plants wither and die. But if we see that event as a causal action on the part of some agent, then that agent is a reaper. So far, so good. But why destroyers and devourers? And what about the impossible cases?

Destruction and devouring are actions in which an entity ceases to exist. The same is true of death. The overall "shape" of the event of death is similar in this respect to the overall "shapes" of the events of destruction and devouring. Moreover, there is a causal aspect to death: the passage of time will eventually result in death. Thus, the overall shape of the event of death has an entity that over time ceases to exist as the result of some cause. Devouring and destruction have the same overall "event-shape". That is, it is the same with respect to causal structure and the persistence of entities over time.

We therefore hypothesized that EVENTS ARE ACTIONS is constrained in the following way: the action must have the same overall event-shape as the event. What is preserved across the mapping is the causal structure, the aspectual structure, and the persistence of entities. We referred to this as "generic-level structure".

The preservation of generic-level structure explained why death is not metaphorized in terms of teaching, or filling the bathtub, or sitting on the sofa. They simply do not have the same causal and overall event structure, that is, they do not share "generic-level structure".

Problem 2: Proverbs

In discussing a collection of Asian figures — proverbs in the form of short poems — the question arose as to what are the limitations on the interpretation of a proverb. Some interpretations are natural; others seem impossible. Why?

Consider the following example from *Asian Figures*, translated by William Merwin (1973).

Blind
blames the ditch

To get some sense of the possible range of interpretations for such a proverb, consider the following application of the proverb:

Suppose a presidential candidate knowingly commits some personal impropriety (though not illegal and not related to political issues) and his candidacy is destroyed by the press's reporting of the impropriety. He blames the press for reporting it, rather than himself for committing it. We think he should have recognized the realities of political press coverage when he chose to commit the impropriety. We express our judgment by saying, "Blind/blames the ditch."

Turner and I observed that the knowledge structure used in comprehending the case of the candidate's impropriety shared certain things with the knowledge structure used in comprehending the literal interpretation of "Blind/blames the ditch". Let us refer to it as the generic-level schema structure our knowledge of the proverb. That generic-level knowledge structure is:

There is a person with an incapacity.
He encounters a situation in which his incapacity in that situation results in a negative consequence.
He blames the situation rather than his own incapacity.
He should have held himself responsible, not the situation.

This is a very general schema characterizing an open-ended class of situations. We can think of it as a variable template that can be filled in in many ways. Here is one way:

The person is the presidential candidate.
His incapacity is his inability to understand the consequences of his personal improprieties.
The context he encounters is his knowingly committing an impropriety and the press's reporting it.
The consequence is having his candidacy dashed.
He blames the press.
We judge him as being foolish for blaming the press instead of himself.

If we view the generic-level schema as mediating between the proverb "Blind/blames the ditch" and the story of the candidate's impropriety, we get the following correspondence:

The blind person corresponds to the presidential candidate.
His blindness corresponds to his inability to understand the consequences of his personal improprieties.
Falling into the ditch corresponds to his committing the impropriety and having it reported.
Being in the ditch corresponds to being out of the running as a candidate.
Blaming the ditch corresponds to blaming the press coverage.

Judging the blind man as foolish for blaming the ditch corresponds to judging the candidate as foolish for blaming the press coverage.

This correspondence defines the metaphorical interpretation of the proverb as applied to the candidate's impropriety. Moreover, the class of possible ways of filling in the generic-level schema of the proverb corresponds to the class of possible interpretations of the proverb. Thus, we can explain why "Blind/blames the ditch" does not mean "I took a bath" or "My aunt is sitting on the sofa" or any of the myriad of things the proverb cannot mean.

All of the proverbs that Turner and I studied turned out to involve this sort of generic-level schema. And the kinds of things that turned up in such schemas seemed to be pretty much the same in case after case. They include:

Causal structure.
Temporal structure.
Event shape; that is, instantaneous or repeated, completed or open-ended, single or repeating, having fixed stages or not, preserving the existence of entities or not, and so on.
Purpose structure.
Modal structure.
Linear scales.

This is not an exhaustive list. But what it includes are most of the major elements of generic-level structure that we discovered. What is striking to us about this list is that everything on it is, under the Invariance Hypothesis, an aspect of image–schematic structure. In short:

If the Invariance Hypothesis is correct, the way to arrive at a general-level schema for some knowledge structure is to extract its image–schematic structure.

The metaphoric interpretation of such discourse forms as proverbs, fables, allegories, and so on seems to depend on our ability to extract generic-level structure. Turner and I have called the process of extracting the generic-level structure from a specific knowledge structure the GENERIC IS SPECIFIC metaphor. We see it as a general mechanism for understanding the general in terms of the specific.

If the Invariance Hypothesis is correct, then the GENERIC IS SPECIFIC metaphor is a minimal metaphor that maps what the Invariance Hypothesis requires it to and nothing more. Should it turn out to be the case that generic-level structure is exactly image–schematic structure, then the Invariance Hypothesis would have enormous explanatory value. It would obviate the need for a separate characterization of generic-level

structure. Instead, it would itself characterize generic-level structure — explaining possible personifications and the possible interpretations for proverbs.

6. The status of the Invariance Hypothesis

The Invariance Hypothesis is just that — an empirical hypothesis. Moreover, its status is anything but clear. In the first place, it is vague in certain respects, since a precise formulation would require knowledge of the full inventory of image–schemas. Secondly, it could be stated in either a strong or a weak form. There are various possible weak forms. For example, one might claim that not all abstract inferential structure is image–schematic, but that only some specified portion of it is. Or one might consider the possibility that image–schematic structure is only one of a number of aspects of generic-level structure.

Of course, the most interesting form of the Invariance Hypothesis is its strongest form:

All metaphorical mappings are partial. What is mapped preserves image–schematic structure, though not all image–schematic structure need be mapped. Furthermore, all forms of abstract inference, all details of image-mappings, and all generic-level structure arise via the Invariance Hypothesis.

At present, there is certainly not sufficient evidence to support this form of the hypothesis. But, as a research strategy, keeping to as strong a form of the Invariance Hypothesis as possible is a good way to investigate just what the limits of invariance are.

7. Conclusion

We began with a discussion of the cognitive and generalization commitments, and I would like to return to that discussion. The Invariance Hypothesis has resulted from an attempt to meet the generalization commitment as well as possible. Even in less than its strongest form, the Invariance Hypothesis would still constitute a grand generalization, a generalization that extends over inference patterns that metaphors preserve, over constraints on image-metaphors, and over aspects of generic-level structure.

What is less obvious is how the Invariance Hypothesis relates to the cognitive commitment. I take it as part of the cognitive commitment to

characterize what abstract concepts are, how they can be understood, and how abstract reason could have been acquired by human beings. The Invariance Hypothesis could play a major role in such an endeavor. It claims that many abstract concepts arise from metaphorical mappings of spatial concepts and that abstract reason arises via metaphorical mapping when the cognitive topology of image-schemas is preserved by the mapping, which in turn preserves the inferential structure of those spatial concepts.

One of the things we know about the evolution of the brain is that structures that evolved in lower animals are used and elaborated on in higher animals. The Invariance Hypothesis claims that certain mechanisms for the perception of spatial relations that appear to be present in lower animals are used by human beings in abstract reasoning — that aspect of human beings that has traditionally been taken as separating man from the lower animals. But biology has shown us that man is not a completely separate life form; rather, human beings use many biological capacities present in animals that evolved earlier. From the evolutionary point of view, of course, it would not be surprising if human reason were to use and build on mechanisms for representing spatial relations that are present in lower animals. Indeed, the idea that abstract reason makes use of spatial perceptual mechanisms present in lower animals makes much more sense than the idea that reasoning came in all at once with man as a totally separate new cognitive faculty. The idea that abstract reason also evolved just makes more biological sense.

University of California at Berkeley

Note

1. This research was supported in part by grants from the Sloan Foundation and the National Science Foundation (IRI-8703202) to the University of California at Berkeley.

The author is Professor of Linguistics and a faculty member of both the Institute for Cognitive Studies and the International Computer Science Institute at the University of California at Berkeley. He is President of the International Cognitive Linguistics Association and a member of the Governing Board of the Cognitive Science Society.

The following colleagues and students helped with this paper in a variety of ways, from useful comments to allowing me to cite their research: Ken Baldwin, Claudia Brugman, Jane Espenson, Sharon Fischler, Ray Gibbs, Adele Goldberg, Karin Myhre, Eve Sweetser, and Mark Turner.

Poetic Works Cited

Auster, Paul (ed.) 1984. *The Random House Book of Twentieth Century French Poetry*. New York: Random House.

Merwin, W.S. 1973. *Asian Figures*. New York: Atheneum.
Merwin, W.S., and Masson, J. Moussaieff, trs. 1981. *The Peacock's Egg*. San Francisco: North Point Press.
Rothenberg, Jerome (ed.) 1985. *Technicians of the Sacred*. Berkeley and Los Angeles: University of California Press.

References

Gibbs, Raymond W., Jr., and O'Brien, Jennifer E.
1989 Idioms and Mental Imagery: The Metaphorical Motivation for Idiomatic Meaning. Psychology Department, University of California at Santa Cruz. Unpublished ms.

Johnson, Mark
1987 *The Body in the Mind: the Bodily Basis of Meaning, Reason and Imagination*. Chicago: University of Chicago Press.

Lakoff, George
1987 *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*. Chicago: University of Chicago Press.
1989 Philosophical Speculation and Cognitive Science. In *Philosophical Psychology*: 2,1.

—, and Claudia Brugman
1986 Argument Forms in Lexical Semantics. In Nikiforidou et al. (eds.). Proceedings of the Twelfth Annual Meeting of the Berkeley Linguistics: 442–454.

—, and Johnson, Mark
1980 *Metaphors We Live By*. Chicago: University of Chicago Press.

—, and Turner, Mark
1989 *More Than Cool Reason: A Field Guide to Poetic Metaphor*. Chicago: University of Chicago Press.

Langacker, Ronald
1987 *Foundations of Cognitive Grammar: Volume I*. Stanford: Stanford University Press.

Rumelhart, David, and Jay McClelland
1986 *Parallel Distributed Processing*. Cambridge: MIT Press.

Sweetser, Eve
1990 *From Etymology to Pragmatics: The Mind-as-Body Metaphor in Semantic Structure and Semantic Change*. Cambridge: Cambridge University Press.

Talmy, Leonard
1985 Force Dynamics in Language and Thought. In *Papers from the Parasession on Causatives and Agentivity*. Chicago: Chicago Linguistic Society.

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