

**NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:**  
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

CMU-CS-81-115

# **Metaphor: An Inescapable Phenomenon in Natural Language Comprehension**

Alternate (short) Title:

**Metaphor Comprehension**

**Jaime G. Carbonell**  
Carnegie-Mellon University  
Pittsburgh, PA 15213

4 May 1981

DRAFT # 4

3.6.2  
JAR

## Table of Contents

### **1. Metaphor as a Key to Extensible Semantic Analysis**

- 1.1 An Opening Argument
- 1.2 Recognition vs. Reconstruction - The Central Issue
- 1.3 Steps Towards Codifying Knowledge of Metaphors
- 1.4 A Glimpse into the Process Model
- 1.5 Two Examples Brought to Light
- 1.6 Freezing and Packaging Metaphors
- 1.7 Wrapping Up

### **2. Part II: Metaphors, Inference and Memory Organization**

- 2.1 Two Metaphors are Better Than One
- 2.2 Why Metaphor?
- 2.3 How Metaphors Structure Inference Processes
- 2.4 Analogical Mappings in Problem Solving
- 2.5 Exploiting the Invariance Hierarchy
- 2.6 Implications for Memory Organization

# Metaphor: An Inescapable Phenomenon in Natural Language Comprehension

Jaime G. Carbonell  
Carnegie-Mellon University  
Pittsburgh, PA 15213

## Abstract

Interpreting metaphors is an integral and inescapable process in human understanding of natural language. Part I of this paper discusses a method of analyzing metaphors based on the existence of a small number of generalized metaphor mappings. Each generalized metaphor contains a *recognition network*, a *basic mapping*, *additional transfer mappings*, and an *implicit intention component*. It is argued that this method reduces metaphor interpretation from a reconstruction to a recognition task. Steps towards automating certain aspects of language learning are also discussed. Part II analyzes analogical mappings underlying metaphors and implications for inference and memory organization. Regularities have been observed indicating that certain types of conceptual relations are much more apt to remain invariant in analogical mappings than other relations, resulting in an induced invariance hierarchy. The central thesis is that human inference processes are governed by the same analogical mappings manifest as metaphors in language.<sup>1</sup>

---

<sup>1</sup>This research was sponsored in part by the Office of Naval Research (ONR) under grant number N00014-79-C-0661, and in part by the Defense Advanced Research Projects Agency (DOD), Order No. 3597, monitored by the Air Force Avionics Laboratory under Contract F33615-78-C-1551. The views and conclusions contained in this document are those of the author, and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

# 1. Metaphor as a Key to Extensible Semantic Analysis

## 1.1 An Opening Argument

A dream of many computational linguists is to produce a natural language analyzer that tries its best to process language that "almost but not quite" corresponds to the system's grammar, dictionary and semantic knowledge base. In addition, some of us envision a language analyzer that improves its performance with experience. To these ends, I developed the *project and integrate* algorithm, a method of inducing possible meanings of unknown words from context and storing the new information for eventual addition to the dictionary [4]. While useful, this mechanism addresses only one aspect of the larger problem - accruing certain classes of word definitions in the dictionary. In this paper, I focus on the problem of augmenting the power of a semantic knowledge base used for language analysis by means of metaphorical mappings.

The pervasiveness of metaphor in every aspect of human communication has been convincingly demonstrated by Lakoff and Johnson [11], Ortony [15], Hobbs [9] and many others. However, the creation of a process model to encompass metaphor comprehension has not been of central concern.<sup>2</sup> From a computational standpoint, metaphor has been viewed as an obstacle, to be tolerated at best and ignored at worst. For instance, Wilks [21] gives a few rules on how to relax semantic constraints in order for a parser to process a sentence in spite of the metaphorical usage of a particular word. I submit that it is insufficient merely to tolerate a metaphor. Understanding the metaphors used in language often proves to be a crucial process in establishing complete and accurate interpretations of linguistic utterances.

## 1.2 Recognition vs. Reconstruction - The Central Issue

There appear to be a small number of *general* metaphors (on the order of fifty) that pervade commonly spoken English. Many of these were identified and exemplified by Lakoff and Johnson [11]. For instance: *more-is-up*, *less-is-down* and the *conduit metaphor* - Ideas are objects, words are containers, communication consists of putting objects (ideas) into containers (words), sending the

---

<sup>2</sup>Hobbs has made an initial stab at this problem, although his central concern appears to be in characterizing and recognizing metaphors in commonly-encountered utterances.

containers along a conduit (a communications medium, such as speech, telephone lines, newspapers, letters), whereupon the recipient at the other end of the conduit unpackages the objects from their containers (extracts the ideas from the words) [17]. Both of these metaphors apply in the examples discussed below. Clearly, not all metaphors are instances of the identified set of general metaphors, but a sampling from narrative and newspaper text indicates that the majority of metaphors are indeed instantiated versions of a few general metaphors. The computational model below applies only to this class of common metaphors. Other, more "creative" metaphors are discussed in part II of this paper.

The computational significance of the existence of a small set of general metaphors underlies the reasons for my current investigation: The problem of understanding a large class of metaphors may be reduced from a reconstruction to a recognition task. That is, the identification of a metaphorical passage as an *instance* of one of the general metaphorical mappings is a much more tractable process than reconstructing the conceptual framework from the bottom up each time a new metaphor-instance is encountered. Each of the general metaphors contains not only mappings of the form: "**X is used to mean Y in context Z**", but inference rules to enrich the understanding process by taking advantage of the reasons why the writer may have chosen the particular metaphor (rather than a different metaphor or a literal rendition).

### 1.3 Steps Towards Codifying Knowledge of Metaphors

I propose to represent each general metaphor in the following manner:

- A Recognition Network contains the information necessary to decide whether or not a linguistic utterance is an instantiation of the general metaphor. On the first-pass implementation I will use a simple discrimination network.
- The Basic Mapping establishes those features of the literal input that are directly mapped onto a different meaning by the metaphor. Thus, Any upward movement in the *more-is-up* metaphor is mapped into an increase in some directly quantifiable feature of the part of the input that undergoes the upward movement.
- The Implicit-Intention Component encodes the reasons why this metaphor is typically chosen by a writer or speaker. Part of this information becomes an integral portion of the semantic representational of input utterances. For instance, Lakoff identifies many different metaphors for love: *love-is-a-journey*, *love-is-war*, *love-is-madness*, *love-is-a-patient*, *love-is-a-physical-force* (e.g., *gravity*, *magnetism*). Without belaboring the point, a writer chooses one these metaphors, as a function of the ideas he wants to convey to the reader. If the understander is to reconstruct those ideas, he ought to know why the particular metaphor was chosen. This information is precisely that which the metaphor

conveys that is absent from a literal expression of the same concept. (E.g., "John is completely crazy about Mary" vs. "John loves Mary very much". The former implies that John may exhibit impulsive or uncharacteristic behavior, and that his present state of mind may be less permanent than in the latter case. Such information ought to be stored with the *love-is-madness* metaphor unless the understanding system is sufficiently sophisticated to make these inferences by other means.)

- A Transfer Mapping, analogous to Winston's Transfer Frames [22], is a filter that determines which additional parts of the literal input may be mapped onto the conceptual representation, and establishes exactly the transformation that this additional information must undergo. Hence, in "Prices are soaring", we need to use the basic mapping of the *more-is-up* metaphor to understand that prices are increasing, *and* we must use the transfer map of the same metaphor to interpret "soar" (= rising *high* and *fast*) as *large* increases that are happening *fast*. For this metaphor, altitude descriptors map into corresponding quantity descriptors and rate descriptors remain unchanged. This information is part of the transfer mapping. In general, the default assumption is that all descriptors remain unchanged unless specified otherwise - hence, the frame problem [13] is circumvented.

In essence, general metaphors provide an expectation setting to comprehend an utterance in a top-down fashion, exploiting knowledge of identified, common, culturally-accepted underlying analogies. The computational benefits of such a method have been proven by ELI [18] and its successors [2] for expectation-based literal language analysis. The expectation-based paradigm can now be extended to metaphor comprehension.

## 1.4 A Glimpse into the Process Model

The information encoded in the general metaphors must be brought to bear in the understanding process. Here, I outline the most direct way to extract maximal utility from the general-metaphor information. Perhaps a more subtle process that integrates metaphor information more closely with other conceptual knowledge is required. An attempt to implement this method in the near future will serve as a pragmatic measure of its soundness.

The general process for applying metaphor-mapping knowledge is the following:

1. Attempt to analyze the input utterance in a literal, conventional fashion. If this fails, and the failure is caused by a semantic case-constraint violation, go to the next step. (Otherwise, the failure is probably not due to the presence of a metaphor.)
2. Apply the recognition networks of the generalized metaphors. If one succeeds, then retrieve all the information stored with that metaphorical mapping and go on to the next step. (Otherwise, we have an unknown metaphor or a different failure in the original semantic interpretation. Store this case for future evaluation by the system builder.)

3. Use the basic mapping to establish the semantic framework of the input utterance.
4. Use the transfer mapping to fill the slots of the meaning framework with the entities in the input, transforming them as specified in the transfer map. If any inconsistencies arise in the meaning framework, either the wrong metaphor was chosen, or there is a second metaphor in the input (or the input is meaningless).
5. Integrate into the semantic framework any additional information found in the implicit-intention component that does not contradict existing information.
6. Remember this instantiation of the general metaphor within the scope of the present dialog (or text). It is likely that the same metaphor will be used again with the same transfer mappings present, but with additional information conveyed. (Often one participant in a dialog "picks up" the metaphors used by the other participant. Moreover, some metaphors can serve to structure an entire conversation.)

## 1.5 Two Examples Brought to Light

Let us see how to apply the metaphor interpretation method to some newspaper headlines that rely on complex metaphors. Consider the following example from the New York Times:

**Speculators brace for a crash in the soaring gold market.**

Can gold soar? Can a market soar? Certainly not by any literal interpretation. A language interpreter could initiate a complex heuristic search (or simply an exhaustive search) to determine the most likely ways that "soaring" could modify gold or gold markets. For instance, one can conceive of a spreading-activation search starting from the semantic network nodes for "gold market" and "soar" (assuming such nodes exist in the memory) to determine the minimal-path intersections, much like Quillian originally proposed [16]. However, this mindless intersection search is not only extremely inefficient, but will invariably yield wrong answers. (E.g., a gold market ISA market, and a market can sell fireworks that soar through the sky - to suggest a totally spurious connection.) A system absolutely requires knowledge of the mappings in the *more-is-up* metaphor to establish the appropriate and only the appropriate connection.

In contrast to an intersection search, consider an application of the general mechanism described in the previous section to the "soaring gold market" example. Upon realizing that a literal interpretation fails, the system can take the most salient semantic features of "soaring" and "gold markets" and apply them to the recognition networks of the general metaphors. Thus, "upward movement" from soaring matches "up" in the *more-is-up* metaphor, while "increase in value or volume" of "gold markets" matches the "more" side of the metaphor. The recognition of our



example as an instance of the general *more-is-up* metaphor establishes its basic meaning. It is crucial to note that without knowledge that the concept *up* (or ascents) may map to *more* (or increases), there appears to be no general tractable mechanism for semantic interpretation of our example.

The transfer map embellishes the original semantic framework of a gold market whose value is increasing. Namely, "soaring" establishes that the increase is rapid and not firmly supported. (A soaring object may come tumbling down implies that rapid increases in value may be followed by equally rapid decreases). Some inferences that are true of things that soar can also transfer: If a soaring object tumbles it may undergo a significant negative state change, therefore the gold market (and those who ride it) may suffer significant negative state changes. However, physical states map onto financial states.

The *less-is-down* half of the metaphor is, of course, also useful in this example, as we saw in the preceding discussion. Moreover, this half of the metaphor is crucial to understand the phrase "bracing for a crash". This phrase must pass through the transfer map to make sense in the financial gold market world. In fact, it passes through very easily. Recalling that physical states map to financial states, "bracing" maps from "preparing for an expected sudden physical state change" to "preparing for a sudden financial state change". "Crash" refers directly to the cause of the negative physical state change, and it is mapped onto an analogous cause of the financial state change.

*More-is-up, less-is-down* is such a ubiquitous metaphor that there are probably no specific intentions conveyed by the writer in his choice of the metaphor (unlike the *love-is-madness* metaphor). The instantiation of this metaphor should be remembered in interpreting subsequent text. For instance, our "gold market" example continued:

**Analysts expect gold prices to hit bottom soon, but investors may be in for a harrowing roller-coaster ride.**

We require the context of: "up means increases in the gold market, and down means decreases in the same market, which can severely affect investors" before we can hope to understand the "roller-coaster ride" as "unpredictable increases and decreases suffered by speculators and investors".

Consider briefly a second example:

**Press Censorship is a barrier to free communication.**

I have used this example before to illustrate the difficulty in interpreting the meaning of the word "barrier". A barrier is a physical object that disallows physical motion through its location (e.g., "The fallen tree is a barrier to traffic"). Previously I proposed a semantic relaxation method to understand an "information transfer" barrier. However, there is a more elegant solution based on the conduit metaphor. The press is a conduit for communication. (Ideas have been packaged into words in newspaper articles and must now be distributed along the mass media conduit.) A barrier can be interpreted as a physical blockage of this conduit thereby disallowing the dissemination of information as packaged ideas. The benefits of applying the conduit metaphor is that only the original "physical object" meaning of barrier is required by the understanding system. In addition, the retention of the basic meaning of barrier (rather than some vague abstraction thereof) enables a language understander to interpret sentences like "The censorship barriers were lifted by the new regime." Had we relaxed the requirement that a barrier be a physical object, it would be difficult to interpret what it means to "lift" an abstract disallowment entity. On the other hand, the lifting of a physical object implies that its function as a disallower of physical transfer no longer applies; therefore, the conduit is again open, and free communication can proceed.

In both our examples the interpretation of a metaphor to understand one sentence helped considerably in understanding a subsequent sentence that referred to the metaphorical mapping established earlier. Hence, the significance of metaphor interpretation for understanding coherent text or dialog can hardly be overestimated. Metaphors often span several sentences and may structure the entire text around a particular metaphorical mapping (or a more explicit analogy) that helps convey the writer's central theme or idea. A future area of investigation for this writer will focus on the use of metaphors and analogy to root new ideas on old concepts and thereby convey them in a more natural and comprehensible manner. If metaphors and analogies help humans understand new concepts by relating them to existing knowledge, perhaps metaphors and analogies should also be instrumental in computer models that strive to interpret new conceptual information.

## 1.6 Freezing and Packaging Metaphors

We have seen how the recognition of basic general metaphors greatly structures and facilitates the understanding process. However, there are many problems in understanding metaphors and analogies that we have not yet addressed. For instance, we have said little about explicit analogies found in text. I believe the computational process used in understanding analogies to be the same as that used in understanding metaphors. The difference is one of recognition and universality of

acceptance of the underlying mappings. That is, an analogy makes the basic mapping explicit (sometimes the additional transfer maps are also detailed), whereas in a metaphor the mapping must be recognized (or reconstructed) by the understander. However, the general metaphor mappings are already known to the understander - he need only recognize them and instantiate them. Analogical mappings are usually *new* mappings, not necessarily known to the understander. Therefore, such mappings must be spelled out (in establishing the analogy) before they can be used. If a mapping is often used as an analogy it may become an accepted metaphor; the explanatory requirement is suppressed if the speaker believes his listener has become familiar with the mapping.

This suggests one method of learning new metaphors. A mapping abstracted from the interpretation of several analogies can become packaged into a metaphor definition. The corresponding subparts of the analogy will form the transfer map, if they are consistent across the various analogy instances. The recognition network can be formed by noting the specific semantic features whose presence was required each time the analogy was stated and those that were necessarily referred to after the statement of the analogy. The most difficult part to learn is the intentional component. The understander would need to know or have inferred the writer's intentions at the time he expressed the analogy.

Two other issues we have not yet addressed are: 1) Not all metaphors are instantiations of a small set of generalized metaphor mappings. 2) Many metaphors appear to become frozen in the language, either packaged into phrases with fixed meaning (e.g., "prices are going through the roof", an instance of the *more-is-up* metaphor), or more specialized entities than the generalized mappings, but not as specific as fixed phrases. I set the former issue aside remarking that if a small set of general constructs can account for the bulk of a complex phenomenon, then they merit an in-depth investigation. Other metaphors may simply be less-often encountered mappings. The latter issue, however, requires further discussion.

I propose that *typical instantiations* of generalized metaphors be recognized and remembered as part of the metaphor interpretation process. These instantiations will serve to grow a hierarchy of often-encountered metaphorical mappings from the top down. That is, typical specializations of generalized metaphors are stored in a specialization hierarchy (similar to a semantic network, with ISA inheritance pointers to the generalized concept of which they are specializations). These typical instances can in turn spawn more specific instantiations (if encountered with sufficient frequency in the language analysis), and the process can continue until until the fixed-phrase level is reached.

Clearly, growing all possible specializations of a generalized mapping is prohibitive in space, and the vast majority of the specializations thus generated would never be encountered in processing language. The sparseness of typical instantiations is the key to saving space. Only those instantiations of more general metaphors that are repeatedly encountered are assimilated into the hierarchy. Moreover, the number or frequency of required instances before assimilation takes place is a parameter that can be set according to the requirements of the system builder (or user). In this fashion, commonly-encountered metaphors will be recognized and understood much faster than more obscure instantiations of the general metaphors.

It is important to note that creating new instantiations of more general mappings is a much simpler process than generalizing existing concepts. Therefore, this type of specialization-based learning ought to be quite tractable with current technology.

## 1.7 Wrapping Up

The ideas described in this section have not yet been fully implemented in a functioning computer system. I intend to continue incorporating them into the POLITICS parser [6], which is modelled after Riesbeck's rule-based ELI [18].

The philosophy underlying this work is that Computational Linguistics and Artificial Intelligence can take full advantage of - not merely tolerate or circumvent - metaphors used extensively in natural language. In case the reader is still in doubt about the necessity to analyze metaphor as an integral part of any comprehensive natural language system, I point out that that there are over 100 metaphors in the above text, not counting the examples. To illustrate further the ubiquity of metaphor and the difficulty we sometimes have in realizing its presence, I note that each section header thus far and the title of this paper contain undeniable metaphors.

Lest the reader form the mistaken impression that the examples discussed above are in any way "artificial" or "contrived", I offer the following anecdote, leading up to the identification of the "soaring gold markets" example. This author was discussing the pervasiveness and centrality of metaphor with a skeptical colleague [who prefers to remain anonymous]. Voluminous examples of metaphors structuring narratives and dialogs only partially convinced my colleague, who stated that perhaps there was a strong case for investigating metaphor comprehension in fictional narratives and discourse, but *obviously* many useful texts can be fully understood by literal semantic analysis --

"Take the financial pages of the New York Times for example. Now *there* is an instance of a very useful domain where purely literal analysis suffices."

I suggested that we put the matter to a test, conceding that perhaps factual newspaper articles would have fewer metaphors, but metaphors would be there nonetheless. We then went to the departmental lounge, opened the New York Times to the financial pages, and were immediately confronted with a plethora of metaphors. Among them was the "soaring gold market" example. In fact, it appears that the density of metaphors per sentence is significantly higher in most "factual" newspaper accounts than in fictional narratives. [We contemplate potential causes for this counterintuitive, perplexing phenomenon in the following sections.]

A moral that can be drawn from this anecdote is that humans are truly unconscious of the difference between metaphorical and literal passages, apparently comprehending both with equal ease. If such is the case, it presents a potential incongruity with my computational model *vis a vis* human cognitive processing of metaphors. To wit, the model predicts that a measurably greater effort may be necessary to process metaphors than literal passages, in spite of the fact that metaphor processing is simplified considerably by converting a seemingly reconstructive task into a top-down recognition-and-instantiation task. Hence, metaphor-interpretation as a cognitive phenomenon remains wide open to continued investigation. This should come as no surprise, since my process-model rests on the observed regularity that metaphors can be categorized into distinct classes according to the semantic mappings manifest -- a linguistic rather than a psychological observation.

## 2. Part II: Metaphors, Inference and Memory Organization

In addition to language comprehension, metaphors structure memory organization and focus inference processes. Recognition and initial interpretation of metaphors in written or spoken language is only the tip of the proverbial iceberg. Metaphors in narratives, dialogs and informative text are the observable linguistic manifestations of a central, underlying cognitive process. My thesis is that *cognition is dominated by analogical reasoning* [7], in contrast with more rigid reasoning models based on "sunder" principles of formal logic (e.g., deduction, resolution, etc.). Metaphor is the reflection, on the language medium, of analogical thought processes; as such it provides essential clues of the inner functioning of human inference processes. This section discusses the utility of metaphor as a tool to investigate various cognitive processes.

Episodic memory is highly idiosyncratic and structured according to directly-experienced events. Adopting Schank's MOPS structure [20, 19], it is evident that mechanisms are required to structure new information in terms of old, as well as creating new associations in memory (in the process of comprehending new material). Kolodner [10] and Lebowitz [12] consider aspects of grouping knowledge into hierarchical organizations and generalizing episodes. However, the key notion of gaining a better understanding of a new concept by means of structuring the new knowledge in the mold of previously-experienced or well-understood schemas is not investigated. Here is where metaphor can prove an invaluable tool in probing reminding and learning phenomena.

### 2.1 Two Metaphors are Better Than One

Consider the metaphor *inflation is war*, as discussed by Lakoff. Newscasters talk of "*fighting inflation*", "*workers taking a beating from inflation*", "*Carter loosing another round to inflation*", "*inflation overrunning our economy*", "*savings being attacked by inflation*", etc. Of what possible use is this metaphor to the reader (or the writer)? Inflation is an economic phenomenon whose causes, implications, and methods of control are not understood by the public at large<sup>3</sup>. (Indeed, some would say that inflation is poorly understood by politicians, economists and business men alike.) Therefore, the metaphor helps to enrich the knowledge brought to bear in the comprehension

---

<sup>3</sup>The popular adage that inflation is "Too much money *chasing* too few goods" is itself a metaphor, one that suggests different corrective action.

process by transferring corresponding appropriate information from the more familiar adversary-conflict situation. The necessity to enrich and elaborate upon concepts in the understanding process has been amply demonstrated by Bransford and Johnson [3], Anderson [1] and others. The only feasible way to bring knowledge to bear on an ill-understood domain is to construct a metaphor suggesting a useful transfer mapping of factual information and inference rules from a better-understood domain.

Given the general notion that metaphors transfer knowledge from well-understood to more ill-structured domains, three questions arise:

1. How can the transferred knowledge be used? (i.e., what does the metaphor provide that a literal description may lack?)
2. How does the transfer of knowledge actually take place? (i.e., what would constitute a computationally-viable mechanism?)
3. What implications does the utility and pervasiveness of metaphor have on cognitive processes such as memory organization, inference, and learning?

We consider each question in turn.

## 2.2 Why Metaphor?

The knowledge transferred from a richer domain to a more impoverished one via an analogical mapping, triggered by the use of linguistic metaphor, plays a key role in inference processes. Let us return to the *Inflation is War* metaphor. If a newspaper article opens with this metaphor, the entire text can be organized around it. Equating inflation with a personified "enemy" enables one to draw upon the knowledge organized under "adversary conflict" to suggest courses of action (in terms of the organizing metaphor)<sup>4</sup>. For instance, we can now understand that in order to "*vanquish inflation*", we clearly must: "*formulate a battle plan*", "*marshal our forces*", "*take the initiative*", "*go on the offensive*", and "*make a determined attack on inflation*" in order to "*stamp it out of our society*" and remain on the alert for "*future bouts with inflation*". In short, we must "*whip inflation now (WIN)*" as President Ford said when he "*launched his campaign against inflation*". When the metaphor has been drawn, it is reasonably easy to formulate subgoals based on the better-understood source domain. This is the first step in planning purposive behavior.

---

<sup>4</sup>This realization is due to Lakoff

The inferences one can draw on how to deal with inflation are all structured by the initial metaphor. Different metaphors will yield markedly different sets of inferences. In order to illustrate that there is nothing inherently special about the *inflation is war* metaphor, consider another metaphor used to discuss inflation, encountered frequently in the Spanish press, but easily understood once stated: *Inflation is a disease*. Here, the economy is the patient, inflation is the infecting organism that must be driven out of the patient with the help of the physician (the economist who sets national economic policy).<sup>5</sup> Hence, one can "take the pulse of the economy", "diagnose the cause (always placing the blame on external forces - just as disease organisms are an external cause of illness)", "prescribe treatment", "put the economy on a lean diet", "make the medicine palatable to the poor", "wait for private enterprise to recuperate", "perform radical surgery, cutting swollen budgets", "treat the symptoms while the inflation continues to ravage", and "relieve the pain by subsidizing the price of necessities".

The moral that can be drawn from these two examples is twofold:

1. **Inference and planning are directly structured by the analogical mapping underlying a dominating metaphor.** The first inflation metaphor suggested *tactics* against inflation, whereas the second suggested *cures* for inflation. Therefore, the inference mechanism consists of mapping corresponding solutions from the source to the target domains. Hence, the metaphor must equate two problems, one of which is better understood and therefore suggests inferences and plans presumed applicable in solving the second problem. (See [7] for a detailed discussion of analogical problem solving.)
2. **Solutions to problems generated by metaphors are ONLY useful as heuristic problem-solving advice.** No detailed solutions in *war* or *curing disease* can be applied directly to inflation. How would one "shoot bullets at inflation" or "get inflation to sign the Geneva convention"? Similarly, one cannot "intern the economy in a hospital" or "give it an intravenous penicillin injection". Clearly, the underlying analogy does not extend to the object level. However, the *planning* level provides useful information. Therefore the **intentions of the actors** are preserved in the mapping, as is the **causal structure** of the events, but the **instantiations** of the events themselves are lost in the analogy. This observation accords with Winston's analogy mappings based on

---

<sup>5</sup>One can imagine this author's confusion upon reading in a Spanish newspaper about the "national malady", the proposal to "inoculate workers" by cost-of living adjustments, and a "prescription for the national health". What sort of epidemic was on the loose? However, once the metaphor was understood, the text made perfect sense. Since this metaphor *is the way in which inflation is always discussed*, there appeared to be no need to introduce it explicitly. Moreover, no one would admit that inflation was not a disease, as the metaphor so deeply permeated discussions of inflation that metaphorical terms were not recognized as such. In a conversation with local person, the following statement was made in reaction to my statement suggesting inflation was being discussed in terms of a medical metaphor: "Of course, our economy is sick and must be cured, literally! I mean just what I said." This episode should help us step outside our own metaphor and realize that no one can literally *battle* inflation, but that the metaphor is so ingrained in our thinking that we can draw inferences and make statements easily *only if* we rely on the accepted metaphor to structure our reasoning processes.



preservation of causal structure [23] and Gentner's discussion of analogy in scientific theory [8].

## 2.3 How Metaphors Structure Inference Processes

As we discussed in the preceding section, metaphors can establish an expectation setting for comprehending large portions of text. This expectation setting is generated by transferring inferences from the source to the target domain, including the goals and plans that actors in the target domain are expected to pursue. (It is important to realize that the goals of "defeating" or "curing" inflation come from the respective metaphors -- not from the concept of inflation itself. Therefore a language understanding system must tap the metaphor to comprehend exactly what problems are caused by inflation, and what their respective solution strategies ought to be.)

Let us define *creative metaphors* to be the linguistic realizations of large-scale analogical mappings that can structure entire planning episodes. Creative metaphors include the two inflation examples above, Gentner's scientific-theory metaphors [8], and each of the roughly 50 generalized metaphors discussed in [5]. Non-creative metaphors are frozen metaphor instances with fixed meanings, or figures of speech (such as "kick the bucket") whose metaphorical roots can only be traced through their etymology. Non-creative metaphors do not map inferences, as their source domain has been lost in their history, and therefore is not available to the understander. In short, if a metaphor enables one to bring knowledge to bear from an existing domain to a new, less understood domain, we define it as a creative metaphor. The discussion below centers on the process that brings knowledge to bear in understanding new domains.

In order to extract information from an existing domain to a new domain via a metaphor, it is crucial to know what aspects of the existing domain should remain invariant in the mapping, which should be transformed, and which should be ignored. As we saw in the previous section, objects are seldom, if ever, preserved in a metaphorical mapping, whereas planning structures are mapped invariant to the new domain -- in fact it is precisely because planning structure and inferences can be preserved by analogical mappings that metaphors are powerful means of helping an understander formulate reasonable behavior in uncharted domains.

An analysis of some two hundred creative metaphors yields the following empirical observation: There is a well-defined invariance hierarchy among the aspects of a situation that are mapped by a

metaphor. This perceived regularity is remarkably consistent across metaphors in different domains. In fact, metaphors that are rated as "bad metaphors" often violate the invariance hierarchy presented below.<sup>6</sup> Hence, a plausible hypothesis is in that people *expect* certain aspects of the source domain to remain invariant and other aspects to be coerced into corresponding entities in the new domain. These expectations can focus the search for metaphor interpretations. The regularities observed over a large number of metaphors are summarized by the normative invariance hierarchy presented below. The conceptual relations in the hierarchy are listed in decreasing order of expected invariance:

- *A goal-expectation setting for the animate actors involved (if any).* -- Goals, if present in the source domain, are almost always mapped invariant into corresponding entities in the target domain. If the source domain contains animate actors and the target does not, then the goals of the actors will be attributed to the corresponding personified entities in the target domain. E.g., inflation becomes an anthropomorphized malevolent agent in *inflation is war*; therefore the *the goals* of a nation at war are mapped invariant in that inflation must be *fought and defeated*.
- *Planning and counterplanning strategies among competing or cooperating actors.* -- These strategies, almost always preserved intact by an analogy, provide a priority ordering among the goals and suggests possible means for pursuing each goal. Often, the most useful aspect of a metaphor is to enable purposive planning in what previously was too ill-structured a domain.
- *Causal Structures.* -- When the causal structure of the source domain is explicit, it will typically be preserved by the mapping. E.g., medicine cures disease; therefore economic measures will "cure" inflation. In Reddy's conduit metaphor for how people talk of language [17], causal structure abounds. E.g., a blocked conduit prevents physical transfer; therefore press censorship will also block dissemination of ideas.
- *Functional Attributes.* -- The function to which an object in the source domain is typically applied will often be coerced onto an analogous function for the corresponding object in the target domain.
- *Temporal Orderings.* -- Normative planning sequences in the source domain map into potentially applicable planning sequences (instantiated differently) in the target domain, often preserving temporal relations.
- *Natural tendencies.* -- In the celebrated analogy between electric circuits and and hydraulic systems (used to explain Ohm's Law), water "tends" to go down hill, therefore electricity "tends" to go towards the voltage "*drop*". Moreover, thin pipes resist the flow of water, therefore thin wires "resist" the "*flow*" of electricity.<sup>7</sup>

---

<sup>6</sup>This observation is based on data collected by Lakoff, Gentner, Ortony and others.

<sup>7</sup>Electricity is actually not a flow of electrons, but we always think of it that way because the hydraulics metaphor pervades our discussions of electrical phenomena.

- *Social roles.* -- Social relations are sometimes preserved and sometimes not. In a battle there are generals and foot-soldiers; therefore, the war against inflation must be fought by many wage-earners (soldiers) under the direction of economic planners (generals). Since doctors cure the disease directly, the economic planner must shoulder the entire burden, and wage-earners (patients) are essentially powerless with respect to taking an active role in the cure. Both mappings preserve the inferences associated with the social roles in the source domain. However, the more specific roles of "spy" and "submarine commander", are not preserved by the *inflation is war* metaphor.
- *Structural relations.* -- Occasionally structural relations remain invariant in an analogy, but often they are transformed or suppressed. For instance, in the Rutherford solar-system model of the atom, physical relations between the electrons (planets) and nucleus (sun) are remain invariant. (In both case there is an orbit relation as a function of an inverse-square centripetal force). However, Saying "John is at the head of his class" does not preserve the physical structure normally found between a body and a head (the latter being connected to and nourished by the former).
- *Descriptive Properties.* -- These are the least likely properties to be preserved in a metaphor. Wires and pipes are both long and narrow (in the hydraulics metaphor). However, Generals are military men, economic planners are usually academics. The sun is yellow-orange, very large and has sunspots, but as Gentner points out, none of these monadic descriptors apply to the nucleus of the Rutherford atom.
- *Object identity.* -- Objects in the source domain are almost never mapped onto objects of identical type in the target domain. Therefore, there are no tanks, bullets, M16's, attack submarines, uniforms, or field hospitals in the battle against inflation.

## 2.4 Analogical Mappings in Problem Solving

Our discussion suggests that metaphors are a useful means of indexing mappings between goals, planning structures, causal connections, tendencies, relations, and descriptions (in decreasing order of invariance and significance). Not all components are present in every metaphor. The preferred-invariance ordering helps us understand how metaphors may be used to facilitate reasoning processes in new domains, namely:<sup>8</sup>

1. Establish the invariant components of the mapping
2. Establish initial correspondences among the entities in the source and target domains. (This is a very partial correspondence -- only entities that are referenced by an invariant component in the explicit mapping can be directly related.)
3. Goal-correspondence identifies the problems that must be solved in the target domain. [What should one do about inflation? The *disease* metaphor states that it

---

<sup>8</sup>Here I adopt Newell and Simon's Means-Ends Analysis framework for problem solving [14]. The reader is referred to [7] for a more detailed discussion of analogical problem solving.

should be eradicated. The *war* metaphor suggests subjugating it. A comparison of inflation with an overindulging *gourmand* would yield the goal of trimming it down and controlling its scope, but not eliminating it.] Therefore, metaphor actually determines the goals that one ought to pursue in the target domain. Without knowledge of goals little purposive action can take place (i.e., problem solving becomes meaningless, as there are no goal states in the problem space.)

4. Planning strategies invariant under a metaphorical mapping transfer operators from the source to the target domain, hence establishing a problem space [14] and suggesting potentially troublesome interactions among operator preconditions. The *inflation is a disease* metaphor suggests that since administering medicine is a useful operator in the medical domain, a correlate operator ought to be useful in the economic domain. Moreover, medicine is usually an unpleasant experience, therefore the inference is made that its economic correlate would be unpleasant as well. Hence, we speak of giving the economy a strong *dose* of anti-inflationary monetary restraints, and making the policy *palatable* to workers.
5. Causal connections classify operators in the target domain by the differences they reduce (analogized from the source domain). "The pressure of the water is determined by the product of the rate of flow and the cross-section of the pipe" suggests that in order to know the voltage, one can measure the current and resistance. Therefore a way of reducing the KNOW-V goal is to apply the multiply(I, R) operator, reducing the KNOW-V goal to the subgoals KNOW-R and KNOW-I.
6. Natural tendencies, social roles and structural relations provide information about the applicability conditions of operators [E.g., who can administer medicine (decide economic policy)?], and provide heuristic guidance to planning processes in the new domain. [E.g., Wars are costly and people must make personal sacrifices; therefore in battling inflation the cost should be taken into account, and the planner should be ware of potential problems caused by those who are unwillingly called upon to make the sacrifices.]
7. Temporal-progressions suggest macro-operators (typically useful sequences of operators). In treating a disease we first must identify the cause, then prescribe medicine then wait patiently for it to take effect. In war we marshal our forces (no searching for a suitable the enemy is needed, as the enemy is known at the start of hostilities), then attack (no waiting for the attack to take effect is necessary). Therefore we see two very different general plans suggested by the two metaphors. However, recall that the metaphors shared the same general goal. It is typically the case that most metaphors used to explain a particular ill-defined situation will share common goals and diverge increasingly as one traverses down the invariance hierarchy.

Reiterating the central theme of this section: metaphors provide a problem space, including a goal state, operators indexed by differences they may reduce, and normative plans that may prove useful. In essence, they make problem solving possible in what may previously have been too ill-structured a situation to make any progress. Metaphors do *not*, however, provide any *canned* solutions applicable directly to new problems -- such would be an unreasonable expectation.

## 2.5 Exploiting the Invariance Hierarchy

The invariance hierarchy provides a first-pass solution to an apparently simple phenomenon that had perplexed some investigators, including this writer. When we hear that "John is a fox" we interpret it to mean "John is sly", not "John has pointed ears and a bushy tail." Similarly, we interpret "John is a pig" as a remark on his personal habits or his obesity, rather than a statement that John lives in a farm and has a curly tail. A partial answer to this problem lies in knowing the most salient feature of the animal to whom a person is compared. However, a more complete answer is provided by exploiting the invariance hierarchy in the following manner: Consider the animal (source domain) and scan down the hierarchy stopping at the first entry for which we have a commonly known fact. For foxes we stop at planning/counterplanning -- folk wisdom tells us that foxes are very adept at devious counterplanning behavior. Hence, we never reach the physical descriptors of a fox. For pigs we may stop at either "natural tendency" (if we believe that pigs tend to get fat) or at "social role" -- folk wisdom asserts that pigs play a distinct role in the animal kingdom as the least hygienical of all animals. If we heard "John is a Giraffe", we find no common knowledge anywhere in the hierarchy until we reach physical attributes. Here we pick the most salient ones (e.g., height and/or length of neck), to understand the metaphor. The key to the process is that comparisons along the higher-invariance entries in the hierarchy are preferred. Once a high invariant property is found, no lower ones are considered. This is crucial to understand "John is an elephant" as a remark on the length of his memory (or his capacity for work), not the length of his nose (trunk), although the latter is perhaps the single most salient feature of elephants. Physical descriptors, however, are ranked low in the hierarchy.

## 2.6 Implications for Memory Organization

We have outlined how reasoning based on metaphors may proceed. Now, consider another aspect of metaphorical reasoning: *How are metaphors formed in the first place?* Given the ubiquity of metaphor, it becomes strikingly apparent that humans generate metaphors as readily as they understand them, occasionally unconscious of the fact that they are creating (or more often instantiating) metaphors. The question that must be posed is more specific: *What memory organization could enable, facilitate and encourage the continuous creation of metaphors?*

If we assume that the invariance hierarchy is roughly correct, it provides a best-first criterion for searching a content-addressed episodic memory, organized along the general lines of Schank's

MOPS [20, 19]. In investigating reminding and inference phenomena, Schank asserts that detecting similarities at every level of abstraction is the key to human memory organization. Accepting this notion requires one to have a means of computing similarities among large numbers of potentially-relevant episodic traces, both for memory access and update. The hierarchy above suggests that goal similarities are crucial, planning-level similarities are almost as important, and similarities across other dimensions are of progressively lower importance. Hence, if memory were organized according to the computational criteria required for metaphor comprehension, it follows that a hierarchical structure would result, where the categories formed are largely determined by groupings along the entries in the invariance hierarchy, the more invariant entries corresponding to more global organizing categories. The actual content of the hierarchical memory is determined primarily by the idiosyncratic experience of the individual. Therefore, memory searches for "good metaphors" (those preserving high-invariance properties) require less work (either to generate or comprehend) and may prove more rewarding for the understander as they index relevant memory more readily.

Metaphor is a linguistic realization of an inference phenomenon. As such, it should reflect underlying memory structure, as well as suggest the types of inferences people can perform most readily. If we ask *why* creative metaphors are used, the most logical answer appears to be that the writer is trying to induce the reader to perform the necessary inferences required to comprehend the new material. Metaphor serves as a vehicle to suggest a fruitful domain from which the relevant inferences can be mapped onto the new domain. Hence, when Senator Joe McCarthy referred to Communism as a "dreaded plague", he was inducing, in the minds of his listeners, the inference that communism must be actively "eradicated" or it will spread. The metaphor is effective only because the appropriate inference structure was already in existence in the source domain, and McCarthy knew this at the time he created the metaphor.

An interesting avenue of future research is automating metaphor generation. If the model discussed here is essentially correct, metaphor generation requires that the writer have a model of the knowledge state (including goals, strategies, beliefs, etc.) of his reader, as well as an integrated episodic memory where the computed similarity metric incorporates the invariance hierarchy. (I.e., two domains are considered similar if the same types of problems and inference processes are present in planning effective behavior in both domains.)

In order to clear possible misconceptions, I emphasize that no distinct, localized, "conscious" existence for the invariance hierarchy is postulated as part of a human memory model. My hypothesis

is that *the regularities manifested in the hierarchy are epiphenomenal reflections of human memory organization and inference mechanisms*. As such, the invariance hierarchy summarizes a phenomenon that must be explained by comprehensive memory-organization models, and hence it ought to be taken into account in the model-formulation process.

## References

1. Anderson, J. R., *Language, Memory and Thought*, New Jersey: Erlbaum, 1976.
2. Birnbaum, L. and Selfridge, M., "Conceptual Analysis in Natural Language," in *Inside Computer Understanding*, R. Schank and C. Riesbeck, eds., New Jersey: Erlbaum Assoc., 1980, pp. 318-353.
3. Bransford, J. D. and Johnson, M. K., "Considerations of Some Problems of Comprehension," in *Visual Information Processing*, W. G. Chase, ed., New York: Academic Press, 1973.
4. Carbonell, J. G., "Towards a Self-Extending Parser," *Proceedings of the 17th Meeting of the Association for Computational Linguistics*, 1979, pp. 3-7.
5. Carbonell, J. G., "Metaphor - A Key to Extensible Semantic Analysis," *Proceedings of the 18th Meeting of the Association for Computational Linguistics*, 1980.
6. Carbonell, J. G., "POLITICS: An Experiment in Subjective Understanding and Integrated Reasoning," in *Inside Computer Understanding: Five Programs Plus Miniatures*, R. C. Schank and C. K. Riesbeck, eds., New Jersey: Erlbaum, 1981.
7. Carbonell, J. G., "Learning by Analogy: Skill Acquisition in Reactive Environments," in *Machine Learning*, R. S. Michalski, J. G. Carbonell and T. M. Mitchell, eds., Palo Alto, CA: Tioga Pub. Co., 1981.
8. Gentner, D., "The Structure of Analogical Models in Science," Tech. report 4451, Bolt Beranek and Newman, 1980.
9. Hobbs, J. R., "Metaphor, Metaphor Schemata, and Selective Inference," Tech. report 204, SRI International, 1979.
10. Kolodner, J. L., *Retrieval and Organizational Strategies in Conceptual Memory: A Computer Model*, PhD dissertation, Yale University, Nov. 1980.
11. Lakoff, G. and Johnson, M., *Metaphors We Live By*, Chicago University Press, 1980.
12. Lebowitz, M., *Generalization and Memory in an Integrated Understanding System*, PhD dissertation, Yale University, Oct. 1980.
13. McCarthy, J. and Hayes, P. J., "Some Philosophical Problems from Artificial Intelligence," in *Machine Intelligence 6*, Meltzer and Michie, eds., Edinburgh University Press, 1969.
14. Newell, A. and Simon, H. A., *Human Problem Solving*, New Jersey: Prentice-Hall, 1972.
15. Ortony, A., "Metaphor," in *Theoretical Issues in Reading Comprehension*, R. Spiro et al, eds., Hillsdale, NJ: Erlbaum, 1980.
16. Quillian, M. R., "Semantic Memory," in *Semantic Information Processing*, Minsky, M., ed., MIT Press, 1968.
17. Reddy, M., "The Conduit Metaphor," in *Metaphor and Thought*, A. Ortony, ed., Cambridge University Press, 1979.



18. Riesbeck, C. and Schank, R. C., "Comprehension by Computer: Expectation-Based Analysis of Sentences in Context," Tech. report 78, Computer Science Department, Yale University, 1976.
19. Schank, R. C., "Reminding and Memory Organization: An Introduction to MOPS," Tech. report 170, Yale University Comp. Sci. Dept., 1979.
20. Schank, R. C., "Language and Memory," *Cognitive Science*, Vol. 4, No. 3, 1980 , pp. 243-284.
21. Wilks, Y., "Knowledge Structures and Language Boundaries," *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, 1977 , pp. 151-157.
22. Winston, P., "Learning by Creating and Justifying Transfer Frames," Tech. report AIM-520, AI Laboratory, M.I.T., Jan. 1978.
23. Winston, P. H., "Learning and Reasoning by Analogy," *CACM*, Vol. 23, No. 12, 1979 , pp. 689-703.

| REPORT DOCUMENTATION PAGE                                                                                                                             |                       | READ INSTRUCTIONS<br>BEFORE COMPLETING FORM                    |
|-------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|----------------------------------------------------------------|
| 1. REPORT NUMBER<br>CMU-CS-81-115                                                                                                                     | 2. GOVT ACCESSION NO. | 3. RECIPIENT'S CATALOG NUMBER                                  |
| 4. TITLE (and Subtitle)<br>METAPHOR: AN INESCAPABLE PHENOMENON IN NATURAL<br>LANGUAGE COMPREHENSION                                                   |                       | 5. TYPE OF REPORT & PERIOD COVERED<br>Interim                  |
|                                                                                                                                                       |                       | 6. PERFORMING ORG. REPORT NUMBER                               |
| 7. AUTHOR(s)<br>Jaime G. Carbonell                                                                                                                    |                       | 8. CONTRACT OR GRANT NUMBER(s)<br>N00014-79-C-0661             |
| 9. PERFORMING ORGANIZATION NAME AND ADDRESS<br>Carnegie-Mellon University<br>Computer Science Department<br>Pittsburgh, PA. 15213                     |                       | 10. PROGRAM ELEMENT, PROJECT, TASK<br>AREA & WORK UNIT NUMBERS |
| 11. CONTROLLING OFFICE NAME AND ADDRESS                                                                                                               |                       | 12. REPORT DATE<br>May 1981                                    |
|                                                                                                                                                       |                       | 13. NUMBER OF PAGES<br>25                                      |
| 14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)                                                                           |                       | 15. SECURITY CLASS. (of this report)<br>UNCLASSIFIED           |
|                                                                                                                                                       |                       | 15a. DECLASSIFICATION/DOWNGRADING<br>SCHEDULE                  |
| 16. DISTRIBUTION STATEMENT (of this Report)                                                                                                           |                       |                                                                |
| 17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)<br><br>Approved for public release; distribution unlimited |                       |                                                                |
| 18. SUPPLEMENTARY NOTES                                                                                                                               |                       |                                                                |
| 19. KEY WORDS (Continue on reverse side if necessary and identify by block number)                                                                    |                       |                                                                |
| 20. ABSTRACT (Continue on reverse side if necessary and identify by block number)                                                                     |                       |                                                                |