Metaphor - A Key to Extensible Semantic Analysis

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

Abstract

Interpreting metaphors is an integral and inescapable process in human understanding of natural language. 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 the method reduces metaphor interpretation from a reconstruction to a recognition task. Implications towards automating certain aspects of language learning are also discussed.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 [1]. 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 [4], Ortony [6], Hobbs [3] and many others. However, the creation of a process model to encompass metaphor comprehension has not been of central concern.² From a computational standpoint, metaphor has been viewed as an obstacle, to be tolerated at best and ignored at worst. For instance, Wilks [9] 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.

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 [4]. 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 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). Both of these metaphors apply in the examples discussed below.

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 usage 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).

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

¹This research was sponsored 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.

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

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 [10], 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 [5] is circumvented.

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.)

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 [7]. 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 comparison, 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 -> 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 -> 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, had our example continued:

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

We would have needed 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 could 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 disenables 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 disenabling 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 disenabling entity. On the other hand, the lifting of a physical object implies that its function as a disabler 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.

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 in 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: Not all metaphors are instantiations of a small set of generalized metaphor mappings. 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 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.

7. Wrapping Up

The ideas described in this paper have not yet been implemented in a functioning computer system. I hope to start incorporating them into the POLITICS parser [2], which is modelled after Riesbeck's rule-based ELI [8].

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 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 and the title of this paper contain undeniable metaphors.

8. References

1. Carbonell, J. G., "Towards a Self-Extending Parser," *Proceedings of the 17th Meeting of the Association for Computational Linguistics*, 1979, pp. 3-7.
2. 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, 1980.
3. Hobbs, J. R., "Metaphor, Metaphor Schemata, and Selective Inference," Tech. report 204, SRI International, 1979.
4. Lakoff, G. and Johnson, M., *Metaphors We Live By*, Chicago University Press, 1980.
5. McCarthy, J. and Hayes, P. J., "Some Philosophical Problems from Artificial Intelligence," in *Machine Intelligence 6*, Meltzer and Michie, eds., Edinburgh University Press, 1969.
6. Ortony, A., "Metaphor," in *Theoretical Issues in Reading Comprehension*, R. Spiro et al. eds., Hillsdale, NJ: Erlbaum, 1980.
7. Quillian, M. R., "Semantic Memory," in *Semantic Information Processing*, Minsky, M., ed., MIT Press, 1968.
8. 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.

9. Wilks, Y., "Knowledge Structures and Language Boundaries," *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, 1977 , pp. 151-157.
10. Winston, P., "Learning by Creating and Justifying Transfer Frames," Tech. report AIM-520, AI Laboratory, M.I.T., Jan. 1978.