

Understanding New Metaphors

James H. Martin

Berkeley Artificial Intelligence Research
Computer Science Division
University of California
Berkeley, California
94720

Abstract

Metaphors are widespread in natural language. Most natural language processing systems have failed to adequately deal with this issue. A two part approach to this problem involving a Metaphoric Lexicon and a dynamic extension system is described. The extension mechanism operates by analogically mapping previously understood mappings to account for new uses based on known regularities in the English metaphor system. A system based on this approach has been applied to the UNIX domain. Results indicate that this domain exhibits all the predicted regularities and that understanding new senses is possible using these regularities

1. Introduction

Metaphors are widespread in natural language. They are a conventional part of the language and are easily understood through the application of specific knowledge. Metaphors only become a problem when the hearer does not possess knowledge of the metaphor underlying the utterance. Therefore a natural language processing system must have specific knowledge about the conventional metaphors in the language to analyze and generate metaphors. In addition it must have a means of dynamically extending that knowledge when a new use is encountered. In this paper I describe a system that can perform this task. Given a base of knowledge about the conventional metaphors of English and an instance of a usage not covered by that knowledge, it can dynamically discern the new meaning and extend the knowledge base to cover this new use.

The approach is based on several observations. The first is that almost all instances of metaphors are specific instances of a small set of abstract conceptual metaphors (see Lakoff 1980). The second observation is that the specific lexicalized conventional metaphors of English exhibit several systematic regularities. The regularities inherent in the metaphorical system are exploited by the metaphor extension algorithm in an attempt to process new metaphors. The basic assumption is that the new metaphor is conventional and can be expected to fit into the known set of metaphors in a predictable way. The approach is knowledge-based in that it requires knowledge about the conventional metaphors of English, knowledge about the regularities exhibited by those metaphors and knowledge about the conceptual domains that they are applied to.

2. Related Research

Previous approaches have fallen into two distinct classes. The first approach attempts to compute the meaning of the metaphor from knowledge about the literal language of the metaphor and conceptual knowledge of the target domain. The work of Russell (1974), Wilks (1975, 1978), Winston (1979) Carbonell (1981), DeJong and Waltz (1983), and Gentner (1983) exemplify this approach. The other approach attempts to represent knowledge about conventional metaphors. The various structured mapping approaches Carbonell (1981), Jacobs (1985), Norvig (1986), Martin (1986) all assert that the processing of metaphors involves the application of pre-stored knowledge structures representing conventional metaphors. Most of these approaches were inspired by the work of Lakoff and Johnson (1980), which showed that metaphor was a widespread mechanism for structuring the lexicon.

3. Processing New Metaphors

When I use the term New Metaphor I am referring to the situation where an utterance is not understood because of the presence of a metaphor about which the system has no knowledge. (For the moment I draw no distinction between a metaphor which is actually a conventional part of English but about which the system has no knowledge and a metaphor which is actually a novel use). It is the task of the Metaphor Extension System to determine if a metaphor is in fact present and if so to determine its meaning.

The metaphor extension approach relies on the existence of two sources of knowledge to perform its task. The first is an abstraction hierarchy of mappings that represents the set of abstract conventional metaphors in English. When a new use is encountered the system first checks to see if this use might be an instance of an abstract metaphor. If so it then attempts to ascertain the meaning of this new instance.

The second source of knowledge I call the Metaphoric Lexicon. This is a lexicon which explicitly represents the polysemous senses of ordinary English words that have a metaphoric basis. The word senses represented in the Metaphoric Lexicon can be seen as individual lexicalized instantiations of the abstract metaphors described above. The metaphor extension approach relies on the fact that the Metaphoric Lexicon is not a random set of idiomatic phrases. Rather, the lexicon exhibits a small set of regularities. This is the basis for the extension system. The system attempts to relate a new use to one of the known uses via one of these regularities. In the following sections I describe each regularity in the lexicon and show how it can be used for learning.

• This research was sponsored in part by the Defense Advance Research Projects Agency (DOD), Arpa Order No. 4031, Monitored by Naval Electronic System Command under Contract No. N00089C0235 and by a GIE Laboratories Fellowship.

3.1. Metaphorical Frame Semantics

Frame semantics is an approach to lexical semantics proposed by Fillmore (1982). The aspect of this approach that is relevant here is the notion that groups of words are mutually defined with respect to a single semantic frame. The words from a given semantic frame have a well defined set of relationships to one another. Consider Fillmore's example of the commercial transaction frame. The words buy, sell and cost among others are defined with respect to this frame.

Metaphorical Frame Semantics asserts that when multiple terms from the same semantic frame are used metaphorically, they will retain an equivalent relationship to each other as they had in the literal semantic frame. For example if the word *enter* used metaphorically then the metaphorical use of the word *exit* will have a meaning which is in the same relationship (ie. opposite action) to the metaphorical use of *enter* as in the literal meaning. This can be exploited for learning in the following way. When a new use is encountered search for words from the same semantic frame that are used metaphorically in the same way as the current use. If one is found then the base relationship can be used to find or create the new meaning.

3.2. Conceptual Closeness

Conceptual Closeness asserts that concepts that are similar will tend to undergo the same metaphors. Polysemous verb senses are sub-categorized according to the kind of direct object they prefer. The metaphorical meanings of these verbs will tend to apply to objects that are similar. Consider the following examples with the verb *take*.

John *took* a walk with Mary.
John *took* a hike with Mary.
John *took* a cruise with Mary.

Conceptual Closeness refers to the closeness in the categorization of the concepts referred to by the words walk, hike and cruise.

When a new use is encountered one can look to see if any known sense of this verb applies to an object known to be similar to the current object. If so then map the meaning of the known sense to the new sense based on the connection between the objects.

3.3. Metaphorical Coherence

Single concepts tend to be structured by a number of distinct metaphors. Metaphorical Coherence refers to the ways that the various metaphors used to structure a single concept relate to one another. This regularity can be viewed as an analogous version of Conceptual Closeness. With Conceptual Closeness we are interested in the relationships among various closely related target concepts undergoing the same metaphor. With Metaphorical Coherence we are interested in the relationships among various metaphors used to structure one concept. Consider the following examples from Lakoff and Johnson (1980).

Our marriage just *ran out of gas*.
Their relationship is *on the rocks*.
Their marriage is *off the tracks*.

Each of these reflect a structuring of the concept *love* with a particular kind of journey. In particular the notion that the a

failure to make progress in the journey corresponds to problems with the relationship. However distinct sub-types of the journey metaphors are used in each example; they are the car trip, railroad and sea voyage frames.

Metaphorical Coherence can be used to understand new usages in a manner similar to Conceptual Closeness. If you hear a new metaphor look to see if the same target concept is structured via a metaphor that is known to cohere with the currently postulated metaphor. The relationship between the two metaphors and the previously understood mapping is used to create a new mapping for the current target.

4. Algorithm and Implementation

The system was applied to the UNIX domain to test the theory. Consider the following examples:

You can *get into* the editor by typing emacs to the shell.
You can *kill* a process by typing X.
My editor just *died*, how can i save the buffers?
Run your file *through* spell *to fix* the spelling.

The italicized words in these examples are common English words with many polysemous senses. The theory predicts that the meanings of these words in the UNIX domain will be related to their other polysemous senses by one or more of the known regularities. The system was implemented and tested as a component of the UNIX Consultant system (Wilensky 1984). Knowledge about UNIX and the Metaphoric Lexicon were represented using the KODIAK (Wilensky 1986) knowledge representation system - an extended semantic network language.

The learning heuristics described in section three are implemented as a two way breadth first search through the KODIAK network. One search spreads out from the source concept as specified by the given verb. The other starts out from the target concept as indicated by the object of the verb. When these two searches intersect the resulting path is checked against three known path types that represent the three major regularities described above. This algorithm is specifically looking for a path through the network from the source to the target domain across some VIEW representing a known metaphor. A VIEW is a structured KODIAK link that is used to represent non-hierarchical metaphoric mappings. This path is then used to analogically map the meaning of the previous use to the current situation. See Martin (1987) for details on the representation of the Metaphoric Lexicon.

Once a path has been selected a procedure specific to the path type is applied. This procedure produces both the meaning of the new use and new VIEW mappings to represent this new use. The way that the previous mapping is transferred to the new use is dependent on the type of the path. One procedure type is defined for each regularity.

5. An Example

In this section an example of Metaphorical Frame Semantics will be described. The path found will be shown, and the final processing to find the new meaning is described.

Consider the following question given to UC:

How can I *get out* <?/lisp?

The initial state of the knowledge base contains the following information: knowledge of containers, knowledge about lisp

and computer processes in general, and a previously understood mapping defining what getting into lisp means. The metaphor extension system starts a search and finds the path show in Figure 1.

This path corresponds to the type associated with frame semantics. It represents a connection from the source concept `exit-action` to `enter-action` via the frame relation

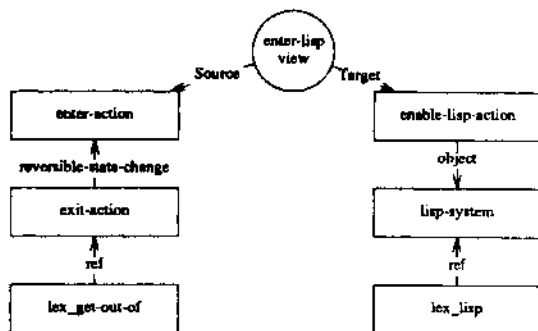


Figure 1: Get out Path Connection

reversible-state-change. The path continues from `enter-action` across the VIEW link to the concept `enable-lisp-use` which is connected directly to the concept `lisp-system` which is the original target concept. This VIEW represents the knowledge about what the phrase getting into lisp means.

The next step is to apply the relation `reversible-state-change` to the concept `enable-lisp-use`. The process of applying a frame relation to a new domain has two possible results. It may simply find a concept already existing in the target domain or it may actually have to create one of the correct type. Here the system finds the initial concept `disable-lisp-use`. This concept is used as the meaning of the new metaphor. Finally a new VIEW is created representing the new use of `exit lisp` to mean `disable`.

6. Conclusions

The key to the success of this system was in the size of initial lexicon. Given a sufficient initial lexicon, knowledge about conventional metaphors, and knowledge about the UNIX domain, the system was able ascertain the meaning of new metaphorical uses. The learning mechanism works basically by analogically mapping previously understood metaphorical mappings. The basis for the success of this approach is the fact the set of conventional metaphors in English is not random. The small set of regularities that it exhibits makes a system based on incremental extension reasonable. Future work is concentrating on issues of representation of the Lexicon and generalization of highly specific mappings to more abstract senses.

References

- Carbonell, J.G. (1980). Metaphor: A key to extensible semantic analysis. In Proceedings of the Third Annual Conference of the Cognitive Science Society (pp. 292-295). Berkeley, CA.
- DeJong, G.F. & Waltz, D.L. (1983). Understanding novel language. Computers and Mathematics with Applications, 131-147.
- Fillmore, C. (1982). Frame semantics. In Linguistics in the Morning Calm (pp. 111-137). Hanshin, Korea: Linguistics Society of Korea.
- Gentner, D. & Stuart, P. (1983). Metaphor as structure mapping: What develops (Technical Report No. 5479). Cambridge, MA: Bolt, Beranek and Newman.
- Hobbs, J. (1979). Metaphor, metaphor schemata, and selective inferencing (Technical Note No. 204). Menlo Park, CA: SRI International.
- Jacobs, P. (1985). A knowledge-based approach to language production. Doctoral Dissertation, Computer Science Division, University of California, Berkeley, CA.
- Lakoff, G. & Johnson, M. (1980). Metaphors We Live By. Chicago, IL: University of Chicago Press.
- Martin, J. (1986). Views from a kill. In Proceedings of the 8th National Conference of the Cognitive Science Society (pp. 728-733). Amherst, MA: Lawrence Erlbaum Associates.
- Martin, J. (1986). Representing and acquiring knowledge about metaphors. In Proceedings of the 3rd Workshop on Theoretical Issues in Conceptual Information Processing (pp. 49-54). Philadelphia, PA.
- Martin, J. (1987). The acquisition of polysemy. To appear in, Proceedings of the Fourth International Workshop on Machine Learning. Irvine, CA: Morgan Kaufmann.
- Norvig, P. (1986). A unified theory of inference for text understanding. Doctoral Dissertation, Computer Science Division, University of California, Berkeley, CA.
- Russell, S.W. (1976). Computer understanding of metaphorically used verbs. American Journal of Computational Linguistics, 44, 45-56.
- Wilensky, R., Arens, Y., & Chin, D. (1984). Talking to UNIX in english: An overview of UC. Communications of the ACM, 27, 574-593.
- Wilensky, R. (1986). Some problems and proposals for knowledge representation (Technical Report No. UCB/CSD 86/294). Berkeley, CA: University of California, Berkeley, Computer Science Division.
- Wilks, Y. (1975). Preference semantics. In E. Keenan, (Ed.), The formal semantics of natural language Cambridge, England: Cambridge University Press.
- Wilks, Y. (1978). Making preferences more active. Artificial Intelligence, 11, 197-223.
- Winston, P. (1978). Learning by creating transfer frames. Artificial Intelligence, 10, 147-172.
- Winston, P. (1980). Learning and reasoning by analogy. Communications of the ACM, 23, 689-703.