Transparently–Motivated Metaphor Generation
(a Ph.D. Proposal)[1]

Technical Report No. 93-14

Mark A. Jones
Department of Computer and Information Sciences
University of Delaware
Newark, Delaware 19716
(Email: *jones@cis.udel.edu*)

[1]This work is supported by Grant #H133E80015 from the National Institute on Disability and Rehabilitation Research. Support has also been provided by The Nemours Foundation.

Abstract

This introduces the phenomenon of transparently-motivated metaphor. It makes a case for its prevalence and reports analysis of some textual goals that such metaphor can achieve. General principles that can explain many uses of such metaphor are discussed. This analysis of metaphors and the resulting principles have led to a methodology for generating transparently-motivated metaphors from general principles. This methodology allows for an account of the conceptual structures behind such expressions, an account of the effects of the metaphors, and the potential for some novel metaphor generation. The issue is approached from the introduction of metaphor generation component. Several details about this component are described, including the forms of the input and output, and how selection rules will be represented, selected and used.

Contents

List of Acronyms:

AP: *Action Predicate*
CV: *Confidence Value*
MT: *Matching Template*
OT: *Output Template*
T-M: *Transparently-Motivated*

1 Introduction

Significant attention has been given to the study of developing a computational system to understand metaphors. The motivation behind such work is straightforward: we use metaphors, and we want computer systems to be able to understand us. There has been almost no attention devoted to the study of developing a computational system to generate metaphors. This lack of study comes from the assumption that there is little benefit in giving a generation system such a capacity. This work contradicts that assumption. Metaphors, including the kind concentrated upon here, are quite abundant in natural discourse. This is no accident. There are two interrelated reasons for the abundance of metaphors. First, people think in metaphorical terms, and, second, metaphors help achieve textual goals. If a Natural Language Generation (NLG) system is to produce truly natural text, it must have the capability to produce metaphors. We will discuss how human conceptualization can be captured to motivate metaphors, and we also will discuss what textual goals can be achieved with such metaphors.

Our understanding of the world influences the way we express ourselves, and so much of what we naturally say is metaphorical. These metaphorical expressions often reflect conceptual models that are the basis for how we understand the world. Mark Johnson [Joh87] has made some interesting observations about the building blocks of thought, most notably that the building blocks of thought are based closely to our bodily experience. Among the building blocks he has described are attraction, blockage, and containment. When describing the purchase of shares of stock one may say, "I took $2,500 out of my money market account and put it into Exxon common stock." The speaker did not literally put money into the stock, but rather bought stock with the money. This metaphor is based on the simplifying concepts that represent investments as containers that can hold money. When we write and talk we automatically use non-literal expressions that reflect our common conceptual groundings. These lead to very natural and easily understood expressions because we (speaker and hearer) share these common conceptual groundings.

Upon analyzing both spoken and written discourse, I find that the use of metaphor (of the kind described herein) is extremely prevalent to the point that many such expressions go unnoticed. Further analysis reveals that metaphor is frequently used as a textual device to help achieve the goals of the speaker. These goals range from the manipulation of focus and perspective, to creating more understandable descriptions that depend on common conceptual groundings.

Typically, however, a natural language generation system only has available to it the literal information that it is to convey. It does not have the conceptual groundings (that underlie many metaphorical statements) that people have. As a result, the generator does not have available to it a common tool that is useful for achieving many textual goals. Because of this, I argue that metaphor is not an added unnecessary feature of a natural language generation system, but is a necessary element. I would like to supply a generator with the necessary conceptual groundings so it can identify additional non-literal ways of expressing what it could only express before literally. This would yield additional options for the system to express itself.

2 Prevalence of Metaphor

When the topic of metaphor is first mentioned, many readers think primarily in terms of blatant metaphors. For example, Cornell Way [Cor91, p. 28] cites a case where a fever is described as "A stubborn and unconquerable flame Creeps in his veins and drinks the streams of life." Although I do not contest that such examples are metaphors, I am interested in the more transparent metaphors. Such metaphors not only occur in friendly and artistic communication, but also in hard, factual domains. The pervasiveness of metaphors in even rather mundane text is supported not only by my analysis of texts, but also the analysis of others. Carbonell [Car82] has found very frequent use of metaphors in the financial pages of the New York Times.

Consider just a few more expressions that often pass without notice. "Savings and Loan bail out" and "Chrysler bail out," although salient at different points in recent history, appear to come from the same conceptual roots. In literal terms, an entity is in trouble, and is getting outside (government) help. The term bail-out that is used to describe this situation is hardly literal, and maps to something more concrete (and thus closer to our conceptual roots). This concrete situation is that of helping to bail out another's boat before it sinks into the water. Other examples include "the stock market is up", "our relationship is on the rocks" and "my career is running out of gas." Appendix A contains more examples of transparently-motivated metaphors that have been encountered in transcript analysis and conversation.

3 Transparently-Motivated Metaphors

One could object that in the previous examples with words such as *up* (in "the stock market is up" example), there is no metaphor. Instead, there are multiple individual word senses. Admittedly, there appears to be a continuum between polysemy and metaphor [And90]. It is an open question whether certain uses of words reflect what should be treated as extra, individual word-meanings or metaphorical meanings that are expansions from some core [NL87]. For example, should we treat the word *big* as having individual senses for language of the physical domain and another for *significant* (as in, "we have a big problem"), or should we treat the word as having one physical sense and treat the less literal meaning as an extension via a metaphorical mapping? We will err on the side of general mechanisms of metaphoric extension and away from the more idiosyncratic solution of positing extra, explicit word senses. For instance, if we use a general metaphoric mechanism for representing less central meanings of the word *big*, we also can explain less central meanings of similar words like *huge*. In both cases, a metaphoric mapping can carry the physical meaning (i.e., large in size) into the abstract realm (i.e., significant). We can speak of a big problem or a huge problem. Although the magnitude of the significance may be different, they have the same relation to each other as they do in the more literal, physical domain. Martin [Mar90] describes this behavior in terms of words being "core-related." The more general approach described in this work can lead to a more parsimonious representation in the lexicon by capturing the deeper generalizations.

Metaphors are non-literal linguistic expressions that have a mapping between two domains. The more literal domain is called the *tenor* domain, the less literal is the *metaphorical* or *vehicle* domain.[2] In the stock example the financial domain is the tenor and the containment domain is the metaphorical domain.

Metaphors have been divided into three basic categories: dead, conventional, and novel. Dead metaphors are often considered frozen phrases. They may have once been novel metaphors, but the tracing of their meaning has been lost. Conventional metaphors are known for generally following common patterns and associations (e.g., more is up, less is down [LJ80], [Lak87]). New metaphors that are not easily related to specific patterns are considered fresh, novel, or creative.

Metaphors have traditionally been defined by their surface form. This definition has recently moved from appealing to than surface form to appealing to conceptual form. Based on the more inclusive definition, a surface form that has never been seen before can be considered a conventional metaphor if it follows a familiar conceptual pattern [Lak93]. With respect to this inclusive definition, this is a study of generating a subset conventional metaphors. This subset is called *Transparently-motivated (T-M) metaphors* [JM92]. We will soon see that T-M metaphors are more narrowly defined with respect to what they can be based upon, and how they are realized.

When people speak in conventional metaphors, they often take advantage of a great deal of very subtle knowledge. It is very likely for a computational system to be ignorant not only of the relevant information, but also that the information *was* relevant. The problem of generating conventional metaphors via computer is a very large problem. The problem of generating T-M metaphors may not be much simpler, but the qualities of T-M metaphors may be more "forgiving" and allow for a more robust system. T-M metaphors will be more robust in such a way that they are insulated from being rendered useless by minor "mistakes" that a computational system (other than a "canned", "hard coded" one) must be expected to make. Yet, we may have to give up little for this valuable simplification, T-M metaphors have the potential to achieve goals that a text planner would request.

The scope of this work will be limited to generating T-M metaphors. These are metaphors that are motivated by general principles and convey a meaning that is primarily not that of the mapping itself. These involve discussing things in more concrete terms or in terms that are conceptually

[2]This terminology is consistent with that of Richards [Ric36].

grounded in common structures of thought. These are restricted to "verb phrase" metaphors. In other words, we are not interested in metaphorical expressions to be used as noun phrase reference. Rather, we are interested in how a metaphorical use of a verb can be explained in a computational fashion.

The prevalence of such metaphor may not at first be apparent. Transparently-motivated metaphors are such a natural form of expression, built of very familiar concepts, that they often do not have the alerting quality that the most flowery metaphors do. Often, it requires several readings of the same page of text before these metaphors can be recognized. The test used to recognize these metaphors is merely to question the literalness of the expression (although the psychological validity of the distinction of literalness is tenuous [Gib84], [Gib89], computationally, this is often accomplished with semantic restrictions) and to identify the two domains.

Four qualities show that the stock example above is transparently-motivated. It is based on the bodily grounding of containment. It conveys the message of *purchase*, which is more than merely the mapping from containment to purchase. Third, this mapping is not itself even highlighted; rather, it is merely used to convey the purchase message. The fourth quality, tenor anchors, will be discussed in section 3.2.

3.1 Metaphor vs. Analogy and Simile

It is important to distinguish metaphor from analogy and simile. A prototypical analogy is: "The sun is to the solar system as a nucleus is to an atom" [Gen83]. For the purpose of this work, the main distinction between metaphor and analogy is what is being expressed. An analogy makes the basic mapping between the tenor domain and the vehicle domain explicit [Car82, p. 421]. The mapping must be consciously recognized by the reader if it is to be understood. This is because the mapping is directly linked to the meaning of the expression. In metaphor the mapping may facilitate the expression of an additional meaning. Similarly, Cornell Way [Cor91, p. 160] shares this distinction, stating that analogy avoids the real question of what facilitates the isomorphism. Instead, it just states the isomorphism.

Furthermore, analogies typically do not employ motivation from common, ubiquitous conceptual groundings. Instead, they often employ novel background mappings. This is why analogies explicitly mark their background mappings. We are not interested in creating new background mappings. We are interested in generating new and familiar expressions based in common, familiar conceptual mappings.

It should be clear that similes do not qualify as transparently motivated for the same reasons. Their mappings are often novel and the meaning is closely tied to the mapping itself. There are expressions that are typically called metaphors which, according to these distinctions, may perhaps be better categorized as analogies or similes, such as "Men are wolves" [Bla62]. Under close inspection, we see that the mapping here is definitely not transparent. Although not stated in the prototypical form (using the form "is like"), the statement has the effect of a simile. Black's example can be rephrased as "men are like wolves." Here the meaning is preserved, perhaps even made more accessible. It appears that Tourangeau and Sternberg would agree to this analysis. They do not define simile in terms of the particular words "is like". Rather they state "When an *overt* nonliteral *comparison* is made, it is usually called a simile" [TS82, p. 204, emphasis added].

3.2 Tenor Anchor

Another counterexample to T-M metaphors is worth discussion. Consider describing the starting of a business in terms of gardening, with expressions such as, "plant the seed in fertile soil" and "give it plenty of water." Such allegorical metaphors are not T-M because they do not have a referent from the tenor domain that remains unchanged by the metaphor [Jon92]. This referent from the tenor can be said to *anchor* the metaphor to the tenor domain, and may be called a *tenor anchor.* The tenor anchor refers to an item that exists in the tenor domain which occurs in the metaphorical expression. For example in the expression, "put money into stock", *stock* is a tenor anchor because it comes from the tenor (financial) domain and not the metaphorical (containment) domain.

The simplest tenor anchors make no sense in the metaphorical domain (i.e., the previous stock example). This does not preclude the possibility of a tenor anchor also being able to refer to items in *both* the tenor and metaphorical domains. In this more complex situation it is important that something, such as the situational context, arbitrate that the tenor anchor be interpreted from the tenor point of view. Later we will consider an example of two parents describing the leaving of a party with the expression "We should pull up stakes." In this example the *We* is the tenor anchor. Yet, without considering the situation, it not is unreasonable to interpret the *We* as referring to the participants in the metaphorical domain (i.e., that the *We* are people in a situation (such as camping) that may literally pull up stakes). It is the situation that arbitrates this dilemma, and allows *We* to function as a tenor anchor.

An appeal to the situational context does not allow just anything to function as a tenor anchor. A tenor anchor must be a referent in the tenor domain in a literal way. Therefore, one cannot appeal to the situational context to claim that *the seed*, in the allegorical example above, is the tenor anchor. This is because there is no literal referent for *the seed* in the tenor domain (unless the business deals in plant seeds). The appeal to the situational context is only needed when the referent that applies in the tenor domain also may literally apply in the metaphorical domain.

3.3 Summary of the Qualities of T-M Metaphors

Since we have considered several related issues to describe the qualities of T-M metaphors, it may be useful to summarize their qualities in an abbreviated form. Transparently-motivated metaphors:

1. Are based on universal groundings that are often linked to bodily experience.
2. Utilize special lexical expressions that simultaneously convey the (verbal) message and the metaphorical domain.
3. Are subtle in the way that they do not draw attention to themselves as blatant metaphors do. At first glance, these metaphors are often not recognized as non-literal.
4. Retain a referent from the tenor domain.

4 The Uses of Metaphor

Metaphor is not merely a nice added attraction to text whose sole purpose is to make the text more enjoyable by adding a flowery flavor. Instead, metaphor can be harnessed to achieve certain textual goals that are difficult to achieve with strictly literal statements. Some of these goals are discussed here. Examples of transparently-motivated metaphors achieving these goals can be found in appendix A.

4.1 Focus of Attention

Metaphor can be used to make text more coherent by giving the generator a useful tool in controlling the focus of attention. Proper focus cues contribute to more coherent text. The use of the term *focus* here refers to what has been called local focus or center ([Sid79], [GJW83a], [Sur92]) and can be thought of as the topic of a sentence.

Different elements of a sentence can be raised in focus to achieve purposes of text planners at a higher level. A simple example of manipulating focus of attention in a sentence is the passive construction. In the sentence "John hit the ball.", *John* is naturally in focus because he is the actor and the subject of the sentence. But in "The ball was hit by John" the prominence of *the ball* has been raised. This is an example of an accepted method for raising the focus of some element by using syntactic means to raise the prominence of the word to be in higher focus. In this case, prominence was raised by putting that word into the first position in the sentence.

The traditional view of focus of attention is that it is something accomplished by syntactic means, such as by making the desired focus the subject of the sentence (as was done in the sentence above with the passive construction). However, there is a relationship between semantic types and level of focus.

Metaphor can allow us to manipulate the semantics of an item. As a result, it can be used to manipulate focus. So, if a writer wishes to say that "John understood the idea," and also wished to make the idea more prominent, a possible option would be "John grasped the idea" (or even "The idea was grasped by John"). In the metaphorical case, *idea* is in higher focus because the reader is invited to think of the idea in more concrete terms. The idea is more prominent.

Concreteness has been found to correlate positively with several phenomena [PYM68], including free recall [DB66], [Sto29], [Pai67], recognition memory [Gor61] and short-term memory [BE68]. These studies have found that more concrete things tend to be remembered better than less concrete things. Thus, all other things being equal, concrete items can be said to be "more in focus" than abstract items. Although memory tests do not directly equate to focus at the time of reading, such techniques are as appropriate for the question as we can expect. This is due to the inappropriateness of on-line measures (e.g., eye fixation and reading time measurements) to this issue (see [Sin90, p. 21]). Thus the above studies show that more concrete items are given higher focus than non-concrete concepts.

Not only is concreteness important to focusing, but so is whether or not something is considered animate. One way that the value of animacy can be demonstrated is by noting how animacy can be linked to general prominence in some languages. There are many languages that do not allow the passivation of a sentence when the deep subject is more animate than the deep object. In Navajo the expression equivalent to "The rock was carried by the man" would not be acceptable [Kee85, p. 272]. The reason for this is based in the semantic types of the objects. Presumably, placing an inanimate rock before an animate object (man) in focus is so unacceptable that it has become a part of their grammar.

Animacy is also very key to focus tracking algorithms. Focus tracking algorithms track the prominent items and are based on the idea that in a coherent text, focus flows from one sentence to the next in an orderly manner. Various types of anaphoric expressions are important to focus tracking algorithms. Pronouns, for example, are seen as a mechanism for indicating continued focus on an item. At the same time, focus tracking information is used to determine the referent of an anaphor. To understand the expression, "It was very nice", we must look back in the discourse to determine what the "it" is referring to. In order to determine the referent of "it", we must consider what has previously been focused upon in the discourse.

The algorithms used to determine focus can be quite complex and utilize many data structures. The accepted view of what constitutes the focus of a sentence is based upon the previous discourse, the cospecifiers of anaphors in the current sentence, syntax and thematic roles. While a description of focusing algorithms is outside the scope of this work, it is interesting to note that agenthood is generally considered to be highly indicative of focus. Agenthood is so important to computing focus that many (e.g., [Sur92], [Car87]) have followed Sidner's [Sid79] lead in giving agenthood a very prominent role in their focus tracking algorithms. Similarly, the work on centering [GJW83b], [BFP83] gives agenthood a very prominent role, but it does this through the deep subject. There are two ways metaphor can be used to raise the focus of an item. One is through altering its semantic type to that of a concrete object (if it is not already concrete). The other is to raise the semantic type to that of an animate object via agenthood (if it is not already animate).[3]

Given this evidence, it is both reasonable and intuitive to conclude that the semantic types of words/concepts affect the perceived level of focus attributed to those words/concepts. We will model the effects of semantic types on perceived focus levels using a *focus hierarchy*. Items at the top of the hierarchy, by nature of their semantic features, are more likely to be focused than those at the bottom. Currently, there are three general regions in the focus hierarchy. In order from most prominent to least prominent those regions are: animate, concrete and abstract.

One effect of metaphorical statements can be to alter the perceived semantic type of an object (and therefore potentially raise the perceived focus level). Consider, with this focus hierarchy in mind, the metaphorical statement "AI *is no stranger* to object-oriented paradigms" [Ell91]. Notice that the phrase "is no stranger" has the effect of conceptually personifying the objects involved (i.e., AI and object-oriented paradigms) since it is a phrase that, literally, can only be used with humans (or perhaps other animate objects). Compare the perceived focus level with that in a more literal rendition of the sentence such as, "AI and object-oriented paradigms have previously been incorporated together."

Thus one effect of metaphor is the manipulation of focus of attention by manipulating the semantic types of objects in a sentence. In the traditional view of focus of attention, a word is treated as having a static semantics. However, the use of metaphor can make the semantic type of objects more fluid. By using them with a verb that only applies to humans, for example, the objects are pushed up the focus hierarchy toward the position that humans occupy.

Thus one principle that emerges is:

> **Semantic Focusing Principle.** One way to raise the perceived focus of attention of an item is to use a metaphor in which the main verb's literal sense requires objects higher in the focus hierarchy. The effect is to raise the perceived semantic type of the object toward that position in the focus hierarchy.

There are two ways that metaphor can raise the focus of an element. One way is the semantic shift of the element to a more concrete or animate level. The second way is that when the semantic shift is finally expressed in a sentence, the object may be placed in subject position. Both together (the semantic shift and the positioning) should tend to have a more pronounced focusing effect than either in isolation.

4.2 Conceptual Fit

In the last section we saw how metaphor could be used to alter the perceived semantic type of a concept and thus alter its perceived level of focus. Having metaphorical options can also make

[3] Due to the flexibility of the semantic representation used in this system, agenthood does not necessarily imply animacy, but under the scope of most verbs it will.

it possible to reach a better conceptual fit. The notion of a "better conceptual fit" is difficult to define. To this work it means taking advantage of cognitive structures, that are only available in English though metaphor, that make the meaning of the sentence easier to comprehend.

There are really two aspects to the conceptual fit of a concept. To motivate these, consider the following two metaphorical expressions about an abstract topic.

- "I cannot grasp this idea."
- "If an idea hits you, let me know."

The first element of conceptual fit can been seen in the way that *idea* is being described in the above examples. Ideas are being viewed as tangible objects in the physical realm. Examples like this are not novel. It is quite natural for people to talk about abstract concepts in a more concrete sense; this grounding appears to make the concepts more understandable. Notice that making an abstract item concrete is often achieved by metaphor.[4] This idea is captured in another principle.

> **Concreteness Principle.** People naturally tend to describe things in a more concrete manner, even when the issue at hand is rather abstract. Such behavior generally yields more natural and understandable text.

While the concreteness principle might explain some of what is going on in the above two examples, it is not enough to explain the similarity of those metaphors with others such as the following.

- "I just put $84.79 *into* groceries."
- "Consider an object class *in* the object-oriented approach."
- "Anything can be *crammed into* an object oriented framework – including AI."

These metaphors have something in common. In each case, a concept is being viewed as a container, with (abstract) objects viewed (along more concrete lines) as being put into or taken out of the container. This is consistent with Johnson's [Joh87] work, which emphasizes that the building blocks of thought are based on bodily experience. In particular, he describes "schemas" for these building blocks, including one for containment.

Thus another common basis for metaphors is to view things in line with one of these schemas that are commonly held by a population. This leads us to another principle:

> **Conceptual Roots Principle.** It is natural to talk about things in terms of "basic building blocks" of thought that are commonly shared by a community.

[4]This reason for metaphor looks similar to focus control. Making an abstract object more concrete also raises the focus level. The point is that it *also* can make the concept easier to understand, or **grasp**.

4.3 Brevity

Paivio [Pai79] lists three hypotheses explaining the usefulness of metaphor. One is that, through vivid imagery, metaphor makes its expression more memorable. This relates to focus, and conceptual fit. The other two hypotheses listed are compactness, and inexpressibility. The compactness hypothesis recognizes that a metaphor allows large chunks of information to be transferred from vehicle to tenor. In this way metaphor can achieve the goal of brevity. It is natural and briefer to say that something is ". . . considered an object class *in* the object-oriented approach." [Ell91], than to say that it is "considered an object class *with respect to* the object oriented approach."

4.4 Other Uses

Two other uses for metaphor that will not be considered in this study should also be mentioned. They are manipulation of perspective and filling lexical holes. The inexpressibility hypothesis states that a metaphor enables us to describe experiences for which appropriate words are not available. Lexical holes, however, differ from the goals listed. Lexical holes are not themselves a goal that may yield additional texture to the text; they demand that a metaphor be generated.

While the perspective on an item or event is related to the notion of focus discussed above, they are distinct. Rather than concentrating on which object is focused on, perspective has to do with *how* an object is viewed. A given perspective on an item or event causes certain aspects of that item or event to be highlighted (and not others) [McC89]. The following example manipulates perspective to cast enthusiasts of traditional AI programming and the object-oriented programming approach in an antagonistic light, highlighting their differences: "Rather than suffering from two warring factions, we should aim to combine the best of both camps and overcome the weaknesses of each approach" [Ell91]. This metaphorical method for highlighting particular aspects of an item can often be found in texts designed to persuade the reader.

5 Related Work

Until the last few decades the dominate view of metaphor was the objectivist view, which states that the most basic level of description of reality is that of literal terms and propositions. In this view, a metaphor is considered merely a linguistic trick that is always inferior to some existing literal form. A metaphor might not only be imperfect, but also suspicious in that its power could be used to mislead. It could not be used to facilitate useful rational thought. More recently, the most compelling arguments claim that not only is metaphor not necessarily inferior to more literal forms, but that human thought itself is based in the looser don-discrete principles associated with metaphor [Joh87, p. 66]. These non-objectivists think that it is unlikely that people have only literal meanings and propositions in their heads. Instead, they claim that people *understand* more abstract ideas in (metaphorical) terms of more concrete bodily experiences.

While the motivating view behind this work is sympathetic with the non-objectivist view, out of necessity we must borrow from the objectivist approach that what is to be conveyed is a literal statement. In other words, we do not plan to represent metaphor as an underlying form. Rather, we hope to simulate the organic phenomena of metaphor by representing them with a set of rules and processes that operate on the literal form. Thus metaphor is secondary in that sense. This is not to say that we completely agree with Fass and Wilks [FW83], that metaphor recognition may take place on the basis of semantic mismatch. Rather, our view of metaphor understanding is consistent with that of Lytinen [LKB92] and Martin [Mar90], which views metaphor understanding as taking place in parallel with literal understanding. [5]

Mark Johnson has pointed to the importance of recognizing that people see the world in terms of their common experiences and their physical bodies. Since the physical notion of containment is fundamental to our existence (we must contain air, water and food to survive), he finds containment among the fundamental concepts that naturally occur in a person's mind. Johnson suggests that many inferential patterns arise from our bodily experience of containment [Joh87, p. 39]. Johnson argues that this is so pervasive that reasoning such as that found in formal logic is understood in bodily terms. The metaphorical view of propositions as locations, along with the bodily understanding of containment, motivates the key metaphor that: for a proposition to be true is to be located in a definite bounded space.

From this metaphor three basic principles from formal logic can be viewed in bodily terms. 1) The law of the excluded middle ($P \vee \neg P$) arises from the bodily experience that one either is or is not in a region. 2) Double negation ($\neg\neg P \leftrightarrow P$) follows from knowing that when you negate the proposition of being in a region, you cross over the boundary line. Therefore to "not not" will lead one out of, and back into, the region. 3)The rule of transitivity ($(a > b > c) \rightarrow (a > c)$) also can be mapped to containment. If you are in the small region S, and S is in the big region B, then you are in B. In this way, the schemas that Johnson describes form the basis of our very understanding of reasoning. They may be considered the building blocks of thought. Such building blocks of thought come into play when considering how some metaphors are generated.

Unfortunately, Johnson does not provide anything close to a computational representation of any of his schemas. Describing the building blocks of thought is a tricky enterprise. Often he describes them with pictures. He emphasizes the subtlety of the schemas by stating that the picture should not be confused as explaining the schema, but rather as a useful tool to help convey these shared concepts. When describing compulsion, Johnson states, "It is present in my perception of a jet airplane being forced down the runway, or in my understanding of forces acting on continental plates, or (metaphorically) in my sense of being forced by peer pressure to join the PTA. The

[5] This subject is discussed in greater detail in section 5.1.1.

schema proper is not a concrete image or mental picture; rather, it is a more abstract pattern that can be manifested in rich images, perceptions, and events." [Joh87, p. 2]

5.1 Metaphor Understanding Research

Several researchers have made progress with the problems of metaphor recognition and understanding (c.f., [Car82], [Fas91], [Gen83], [Mar90], [TS82] and [Wei85]). There is a degree of background that the study of metaphor generation can borrow from existing metaphor understanding work.

Since the mapping between domains is so important to metaphor processing, research on analogy should have some promise for the study of metaphor. Analogy research has made progress in answering the question of what is carried from one domain to another by a metaphor. Carbonell's [Car82] hierarchy of invariance grades ten conceptual relations in terms of whether they are maintained across domains. So, for example, the hierarchy specifies that goals are maintained more than causal structures that are maintained more than descriptive properties. Gentner's [Gen83] work is also relevant in this regard. Gentner's main contribution is that analogies often map relations between domains, and rarely map attributes of objects between the domains. An example that follows this principle is, "The sun is to the solar system as a nucleus is to an atom." Here we see that the relations are what is relevant, such as that of objects revolving around each other. Attributes, such as that the sun is hot, are usually not mapped across analogies.

Although these findings have more relevance to the study of analogical reasoning than with metaphor generation, it is at least relevant to this work. It also helps shed some light on the differences between the behavior of the general metaphor problem, and that of T-M metaphors. These findings are relevant to determining the aptness of a potential metaphor. When creating a novel metaphor that is outside the scope of the T-M variety, one must be careful that the system *explicitly* considers the qualities of the properties and elements that we intend to have "travel" across domains. Carbonell's and Gentner's work tell us that not all properties are equally capable to make such a journey. By focusing on transparently-motivated metaphors, which are simpler in behavior than the complete set of metaphors, we hope that such principles as those identified by Carbonell and Gentner can be handled implicitly.

5.1.1 MIDAS

Among the most recent work in metaphor understanding is that of James Martin. He has developed MIDAS [Mar90, Mar88, Mar87b, Mar87a] (Metaphor Interpretation, Denotation and Acquisition System), a system that interprets and learns metaphors. A key feature of this system is its ability to understand metaphors that are somewhat new to the system.

The system is based on specific verb-frame information. Several frames may be associated with a single verb. Some of these frames represent and encode literal senses of the verb. Other frames can represent conventional metaphors (e.g., kill-sports-defeat, the interpretation preferred in "McEnroe killed Connors"). The system relies on the filler constraints to determine which interpretation is best. This system does not a priori prefer literal interpretations to metaphorical ones.

The system does not assume that all metaphors are a priori encoded in its store of verb frames. Instead it can learn new metaphors as they are encountered. The learning of new metaphors is based on previous verb knowledge and relationships between items, rather than analogical source-target reasoning [Mar90, p. 129]. The new metaphors are interpreted through their relationship with expressions that the system is currently familiar with.

MIDAS attempts to learn a new metaphor when it encounters a semantic mismatch: that is, when an expression is encountered that does not fit any of the verb frames (literal or non-literal) contained in its knowledge base. The system attempts to "learn" new metaphors by extending a frame associated with a known metaphor. There are three ways that the system can learn a new metaphor by extending from a known metaphor. In the first method, the new metaphor is understood by using a known metaphor that is judged to be similar to the new use. For example assume that the system knows that to kill a conversation means to terminate a conversation. When it encounters an expression that involves killing a text editor, it can recognize that both conversations and text editors can be terminated, and that this is a plausible connection from which to interpret the new metaphor.

The second method is based on core relatedness. "Intuitively, words are said to be core-related if their meanings make reference to a common set of concepts" [Mar90, p. 38]. Recognizing core relatedness leads to being able to carry whole sets of words to new target domains. A simple example of words that are core-related are *enter* and *exit*. It seems that these words have a relationship of opposition regardless of how literally they are being used. Because of this, the system can determine what entering lisp probably means if it already knows what exiting lisp means.

The third method that Martin describes reflects that the two previous strategies can be used together. Therefore, if the system knows what "exit lisp" means, it can determine what "enter mail" means. The reasoning that leads to the appropriate interpretation of *enter* is based this previous interpretation of *exit*, and that it is core-related to *enter*. The other part of the interpretation is on the object side. The familiar knowledge about the nature of lisp must be matched with that of mail.

The metaphoric associations, along with the concepts that comprise the source and target domains, are represented using the KODIAK [Wil86] knowledge representation system. It is in the tradition of KL-ONE. Some insights embodied in this system may be useful in metaphor generation work, at least as background, such as the choice of knowledge that is captured and the notion of core-relatedness.

A Note on Semantic Tests Early work in metaphor understanding has been motivated by gross parsing errors, those that a semantic component would find unreasonable as a literal interpretation. The example of "My car drinks gasoline" is easily recognized as non-literal because it will fail the semantic test that recognizes that inanimate objects (e.g., cars) cannot drink. Therefore, the reasoning follows that a literal interpretation should be tried first, and if it fails, then a metaphorical interpretation should be attempted.

Recently, this approach has received criticism, since metaphor can be such a natural and involuntary form of communication. Martin and others have noted that an understanding system should parse metaphors at the same time that more literal interpretations are parsed – not only after semantic tests fail.[6] The best interpretation should be adopted whether semantic constraints fit a literal interpretation or not. As evidence for the validity of this model, consider "McEnroe killed Connors." If the first pass merely considered literal interpretations, no mismatch would be found because people can kill people. However, if this interpretation were forced to compete on an equal field (at the same part of processing) with the metaphoric interpretations, the sports-defeat interpretation would be preferred.

5.1.2 Limits to what can be borrowed

Although there is much background that we can borrow from metaphor understanding research, there are fundamental differences between the problems of metaphor understanding and metaphor generation. Metaphor understanding research can assume that the input that is given to it is natural. When working in generation, we must be careful to insure that the output is expressed naturally. We must choose when it is best to use metaphor rather than a more literal expression. Therefore, the goals must be carefully examined and exploited.

The primary goal in metaphor understanding is to recognize that the input is metaphorical, and to translate it to a more literal form that has some internal meaning to the understanding system. This is not the problem of metaphor generation; a generation system already has the internal representation of what it is instructed to convey as naturally as possible. It may use a metaphorical expression to achieve its textual goals. So, the knowledge that we use must be suited for choosing the most apt metaphor to generate expressions in the most effective manner to achieve the textual goals. With a large number of possible metaphorical domains to choose from, we must be especially sensitive to the issues of size and tractability.

[6] Still, MIDAS does employ a form of semantic test to determine when it should attempt to extend a metaphor [Mar90, p. 109].

5.2 Previous Non-Literal Generation Research

Early work on non-literal generation concentrated on stylized sub-languages [Kuk83], [KPG86]. Although this work generated non-literal expressions, it did not capture the conceptual principles behind the metaphors generated. Instead, non-literal expressions were not fundamentally distinguished from more literal expressions. These systems could only handle what had been explicitly encoded in their lexicons. There was no representation of the concepts that caused these metaphors to makes sense. Because of this lack of generalization there was no capacity for generating anything non-literal that was not specifically encoded.

In his work on the KING generator, Jacobs [Jac85] developed the Ace knowledge representation system. The purpose of Jacobs's work was parsimony of representation, not a general approach to metaphor generation. In the introduction, he stated that " . . . the explicit representation of referential and metaphorical relationships among concepts are necessary for the exploitation of abstractions in the representation of knowledge about language." [Jac85, p. 20].

The representation scheme that KING used included VIEWs that could be used to represent alternative routes of expression. The information in these VIEWs could be motivated by metaphor. However, any metaphor that could be generated by the system would first essentially have to be individually encoded.

To the system's credit, the specification for a VIEW needed not be at a particular sentence form, but at a somewhat conceptual level. For example, it encoded that the action *punch* could be expressed as a transfer using *give* with *punch* being the object of the giving. The purpose behind the development of VIEWs was not only to represent metaphorical means of expression, but also more literal expressions. Using VIEWs allowed his system to represent giving and taking as transfer-events. This, in turn, permitted the encoding of knowledge about describing transfer-events without requiring a given event to be classified as giving or taking [Jac85, p. 53].

There does not currently exist a computational theory that can account for metaphor generation from general principles and concepts. Such an approach should allow for novel metaphor generation, account for the conceptual structures behind such expressions, and account for the effects of such manipulation.

6 Input and Output Forms

The input to the system will be of two parts. One part of the input will specify the textual goals that the text planner wants the metaphor to accomplish. The other part is the sentence-length message that is to be conveyed.

Below is an input specification followed by the output that the metaphor generation component will produce. The input can be glossed as: "Convey the purchase of 100 shares of Exxon stock for

$8,000 while achieving the compatible goal of achieving a greater conceptual fit." The output of the system could be generated as "Put $8,000 into 100 shares of Exxon stock."

Input Goal: Conceptual Fit

Input Message:

- **Action Predicate (AP)** : Purchase
  - **Theme:** Stock (Quantity: 100; Company: Exxon)
  - **Instr:** Dollars (Quantity: 8000)

Output Message:

- **AP**: Put-Into
  - **Theme:** Dollars (Quantity: 8000)
  - **Goal:** Stock (Quantity: 100; Company: Exxon)

The output of the system may be identical to the input (when it fails to find a suitable metaphor), or it may differ from the original input by specifying a metaphorical way of expressing the original input message (as above). In either case, the output of this component is appropriate for input to a standard sentence generation component. We will concentrate on instances where the output is metaphorical. We will describe how a metaphorical expression might be found that expresses the original content in response to the previously identified goals.

6.1 Message Representation

The representation of the input and output message of this system are of the same form. We are considering a representation that captures the logical form of a message in terms of a set of semantic roles. We call the choice of semantic roles a *semantic representation.* The roles that have been used in the above example are: *AP, Theme, Instr* and *Goal.*[7] After a short discussion of the system's requirements of a semantic representation, we will discuss the all of the roles that this system will use.

Normally the question of the semantic representation output by a text planner is outside the scope of work such as this because the metaphor generator is a module independent and subsequent to the text planner. However, because the system described here must manipulate the input of the text planner and output something of the same form as the text planner's output, it is necessary to make some stipulations concerning the representation. The representation needs to be reasonable for a text planning system to use and at the same time be something that my algorithms can work

[7]The terms *Quantity* and *Company* are treated as modifiers of their respective roles.

with. Without yet going into the details of the algorithms, we can say that their major function is to map representations that refer to actions, objects and some modifiers from one domain into a similar representation in another domain. The ideal set of roles for this work would be small and yet able to account for all the meanings to be expressed. It should do this while capturing proper distinctions and generalizations.

There are two main reasons why this semantic representation must have a small set of general roles. First, these representations will be used in mapping schemas that map between different domains. The roles must be reasonable to describe a wide variety of predicates. The roles should not be different for every predicate used, but should be general enough to describe a logical form involving *any* predicate in any domain. Second, because the rules will be structured hierarchically, it will be useful for the system to be able to represent generalizations about the roles. A very general selection rule may need to list the roles that its children may manipulate (in a more specific manner).

We need the roles to be general, yet they must not be so general as to be meaningless. Presumably the roles have some semantic import (since the text planner must decide what is playing each role in the basis of meaning to be expressed). At the same time, the roles also have some syntactic import (since the sentence generation component subsequent to the metaphor generation component described here must select a syntactic form appropriate for conveying the represented meaning). The semantics of the roles described here will be driven by the action predicate (verb).

In comparison, Paul Jacobs [Jac85, p. 120] posits many different kinds of arguments. Essentially, he posits a different set of roles for each verb (e.g., hugging_1 (hugger, object)). Such a form may have been adequate for making the points Jacobs was concerned with, but fails to capture generalizations in both syntax and semantics across various predicates. These generalizations are imperative for metaphor work that is concerned with mapping representations between domains. Without generalizations the metaphorical mapping rules would have to be written specifically for each different predicate. Rather, I intend to develop a small set of rules that can be used for a wide variety of instances of predicates and arguments.

Unfortunately, there is no completely satisfying set of thematic roles available. Dowty observes [Dow91, p. 548], "Although many linguists seem to assume that linguistic theory should include a finite (and short) language-universal canon of thematic roles–including the familiar members Agent, Patient, Goal, Source, Theme, Experiencer, Instrumental, etc.–no one that I know of has ever attempted to propose a complete list." Developing the definitive set of such roles is beyond the scope of this work. However, the problem of specifying the roles manipulated by this system cannot be ignored.

There have been attempts to unify thematic roles into a very small set. One such attempt is *Localism*. Localism is a particularly interesting because it appeals to metaphoric reasoning in an attempt to unify thematic roles. Localism reduces semantic and conceptual information to concrete spatial information [Fra92, p. 229]. In essence, it metaphorically imposes the domain of spatial structure onto meaning from any domain. This way, the roles demanded by various domains (other than those of physical space) are rendered superfluous.

The idea of approaching thematic roles by appealing metaphorically to the concrete spatial domain is not found only in localism. The meanings of the common roles *Theme*, *Source* and *Goal* are described in terms of spatial relations. A prototypical definition of these roles is a bullet (*Theme*) moving from a gun (*Source*) (along a path) to a target. The bullet is the *Theme* because it is the thing being acted upon. The gun is the (*Source*) and the target is the (*Goal*).[8]

Thematic roles were not developed primarily with the needs of this work in mind. My purpose is not to capture the deepest semantic qualities (e.g., causality, internal change) of the arguments. Rather, this system requires a general argument structure with respect to the verb. Therefore, the distinctions between Agent and Author (Agent causes the action, Author does not), or between Theme and Patient (Patient is internally changed by the action, Theme is not) are not of great importance here. Such information will be assumed to be derivable from the meaning of the verb.

Since the distinctions between some of the traditional roles are not relevant to my purposes, I can combine some of those roles. *Theme* and what others have called *Patient* will be combined into the role *Theme*. Primarily, this allows the *Theme* to be the thing acted upon (or changed) by the action. The qualities of *Author* and *Experiencer* will be included in the role *Agent*. This will expand the qualities of *Agent* to be non-volitional, and even non-acting depending on the verb.

The loose definitions of these roles should provide the flexibility necessary for extensive coverage with a small set of roles. The same role may have quite different properties when applied to two different verbs. Yet, much relevant commonality will be captured. The flexibility is shown in the example on page 20. The use of *Instr* in the input and *Goal* in the output are used in a flexible manner. In the input, the money used to make the purchase is categorized as the *Instr* (instrument). The money is viewed as the object that is used to do the action (purchase). The stock is the object being acted upon (*Theme*). In the output metaphor, the stock is categorized as the *Goal* (container) in that the money is put into. We will see in section 7.1.1 the selection rule (rule 3) that accounts for the example output.

Here are the roles that will be used in this system. The AP (Action Predicate) is the concept that is realized by the main verb of the sentence. This can be thought of as a super-role, because all of the other roles derive their meaning from the filler of the *AP*. The objects in the sentence play some role with respect to the AP. The role *Agent* is generally the deep subject of the sentence. Intuitively, it is the thing doing the action; it does not necessarily imply a volitional element to the action. *Theme* is the role that is often identified with the direct object of a sentence; it is the thing being acted upon. It may or may not be internally changed by the action. *Instr* is the instrument

[8] A role of *Path* may prove to be useful later, but it is not currently a part of the representation.

used in performing the action of the verb. The meanings of *Source* and *Goal* originate from an object traveling from a source to a goal. The *Source* is where the traveling object comes from. *Goal* is often identified with the indirect object of the sentence. A *Goal* can be thought of as a receiver, which is not to be confused with *Benef*, the beneficiary of the action. The *Benef* is who the action is being done for. *Loc* signifies the location of the event taking place.

It is not difficult to come up with English sentences that do not fit snugly into the set of roles described here. However, I intend to limit my attention to this set of roles to describe how my theory works. The roles, themselves, are not important. They are a vehicle for accessing various argument positions of verbs in a general way. The roles described should be sufficient to handle a large number of instances. The metaphor generation theory itself is in no way dependent on the particular roles, although changing the roles will require a change in the mapping schemas which may affect the elegance and parsimony of the mapping rules.

6.2 Specifying the Goals

In addition to a semantic representation capturing the logical form to be generated, input to the metaphor generator is a specification of the goal(s) a potential metaphor is to achieve. We will now discuss in greater detail the specification of the textual goals.

To specify the goal of brevity can be as simple as specifying on or off. However, there is no reason that a threshold for the desired length could not be included. The goal of raising the focus of an element is also very simple to communicate to the system. The goal part of the input need only specify the desire to raise focus and to specify what is to be raised. The element that is to be raised would need to be one of those items that occur in the second part of the input.

Requesting the goal of cognitive fit also can be as simple as on and off. However, there really is no "off". Cognitive fit is almost automatic, as it often comes along with, and is used by, other goals. Perhaps more complex possibilities such as requesting a specific kind of cognitive fit are possible. This would lead to the problem of having a large range of possible specifications. These specifications would need to be limited to the "vocabulary" of top-down knowledge that is in the system.

6.2.1 Focus Request Specification

We now turn to the problem of formally specifying how focus control requests from the text planner will be made of the metaphor generation component. At first glance one might assume that a sufficient focus specification would be one that requested an item to be focused upon. However, we see that such a specification would not be rich enough to explain some data. Consider the following four expressions that convey the same propositional content.
(1) "John understood the idea."

(2) "John grasped the idea."

(3) "The idea was understood by John."

(4) "The idea was grasped by John."

Of the four, example (1) places *idea* in the lowest prominence. This would reflect no request to place the idea in focus at all. At the other extreme, example (4) places *idea* in highest focus of all of the examples. Example (4) describes the idea in concrete terms and simultaneously places idea in surface-subject position. Examples (2) and (3) use the tools of example (4) independently; (2) uses a semantic shift and (3) uses a syntactic manipulation. It is not clear what the focus level distinction is between (2) and (3) other than to say that they both lie between the extremes of (1) and (4).

It appears that specifying a single focus choice may not be sufficient to account for examples (1)–(4). Although there is some tradition of distinguishing some gradation of focus, present approaches to local focus and centering do not distinguish between (1) and (2). Yet, intuitively, we see that the idea is in a higher prominence or focus in (2) than (1) because it is treated as more concrete. Traditionally, focus requests in NLG systems have been limited to a specification of *the* focus. However such a request does not seem powerful enough to capture the subtle focus shifts exhibited by the examples above. We will need a more powerful manner of specifying focus requests to capture this subtlety in a principled way. We will employ an input specification that captures a gradation of focus potentially involving two elements from the input logical form.

To accommodate a gradation of focus, the specification of focus requests will be an ordered list with at most two elements. The first element in the list is requested to be in high focus and the second in medium focus. The input specifications for the previous examples follow.

(1) (John)[9] "John understood the idea."

(2) (John, idea) "John grasped the idea."

(3) (John, idea)[10] "The idea was understood by John."

(4) (idea) "The idea was grasped by John."

An account with more than one element being considered in focus manipulation provides for the interaction between metaphor's semantic manipulation and traditional syntactic manipulation. Limiting the list to two elements is not a claim about the limit to the nature of focus in general. It is instead a limit on the subtlety of the focus requests that this system the will attempt to account for.

[9]There is no metaphor in (1) and (3), however.

[10]Example (3) is a borderline case that may match *both* focus requests (John, idea) and (idea, John).

7 Method

This section describes in detail the method involved in generating T-M metaphors. Instead of prescribing acceptable metaphors, we wish to prescribe general structures that motivate familiar metaphors as well as less familiar ones. Transparently-motivated metaphors are a promising subclass of metaphors in which to work, because of their relatively simple nature (they are more difficult to misinterpret), and because they carry special requirements that restrict the possible search space from which they can be generated. We have begun preliminary work on specifying the basic structures and methodologies that together can generate good metaphors.

It is hardly a trivial matter to develop a computational approach to generating an appropriate metaphorical expression for a given meaning that we wish to convey. This is even more difficult when we consider that specific goals also will need to be met. Generating metaphors is not like finding square roots; there is no single path of reasoning that can be expected to yield a suitable metaphor. This approach can be considered opportunistic because more than one metaphorical domain may be isolated and tried. Failure will result when no suitable domain can be found.

The processing consists of four steps.

1. The input is compared with various selection rules. Those rules that accommodate the input are identified as candidates.

2. Semantic transformations specified by the selection rules are applied.

3. The candidate selection rules are judged for how well they will accomplish the goal.

4. The best alternative is selected.

This approach is dominated by top-down rules and associated mappings that are triggered by the tenor domain. The basis for these metaphors is the conceptual roots shared by a community described by the conceptual fit principle.

Recall from section 4.2 that conceptual fit is a common motivation for metaphor generation. From this reason, two principles for deriving metaphors are specified: the concreteness and the conceptual roots principles. The method introduced here is intended to help implement this second principle. If a system is to generate metaphors that follow from conceptual roots, those roots must be represented in the system. We will need *selection rules* and related mapping information to capture the conceptual roots by mirroring such common metaphorical behaviors as those pointed to in Lakoff and Johnson's work [LJ80], [Lak87]. Selection rules should relate such familiar patterns as, "try to describe arguments in terms of war, and argument structure in terms of buildings", "describe progress in terms of a vehicle moving toward a goal" and "describe securities in terms of containers for money." These rules, and the mapping structures associated with them, are a major part of the attempt at the formal embodiment of human conceptual groundings.

7.1 Selection Rule Structure

The job of a selection rule is to represent functionally a metaphoric mapping so that metaphors can be generated from it. We need several different pieces of information to accomplish this. In addition to specifying the mappings between tenor and metaphorical domains, a selection rule must specify several pieces of information that will be used to determine (1) whether an instance of using the rule is going to generate a reasonable metaphor and (2) whether the rule can be successfully used to satisfy the goal requested.

The following describes the fields a selection rule can have and what those fields are intended to capture. This is not intended to be an exclusive list; as research proceeds additional fields may be identified as useful. Each selection rule has several fields including, *parent*, *confidence value* (CV), *semantic result* information, and a *mapping schema* that consists of a *matching template* (MT) and an *output template* (OT).

The function of the *parent* field is merely to help with the hierarchical structure of the selection rules. It indicates the parent from which the current rule inherits information. This minimizes redundancy by allowing general information to be represented infrequently at a more general level. There will be more about the hierarchical structure in section 7.9.

The *semantic result* information indicates where in the focus hierarchy the items in the generated metaphor will belong. This is used to reason about the control of focus and the adherence to the concreteness principle. Together, the semantic result information and the original semantics of the input objects can be compared to determine the semantic shift accomplished by the metaphor.

The *confidence value* (CV) is used to distinguish the relative chance of success between different selection rules. The CV can be viewed as a subjective rating of the "guarantee of success" of a selection rule. It is a simple rating from 1 to 10. A rating of 10 represents a selection rule that is judged very likely to yield an expression that conveys the meaning through the mapping intended in the selection rule. A low rating indicates that the rule (although probably very generative and creative) is likely to yield unnatural metaphors. Such selection rules may require presently unprescribed constraints. Below a certain threshold (perhaps 5), a rule should not be used unless some extra constraints are satisfied. Other things being equal, given several selection rules that are compatible with the input, we should use the one with the highest confidence value. Of course, this factor is considered along with whether the candidate rules are judged to achieve the requested goal(s).

Each selection rule has at least one *mapping schema* that specifies the mapping between the input action predicate with related roles and the action predicate and roles of the output. It consists of a *matching template* (MT) and an *output template* (OT). The matching template is a general specification of a class of input forms with which it is compatible. The schema can only be used when the matching template matches the given input. The variables in the matching template have counterparts in the output template. The output template specifies the semantic form of the output metaphor.

7.1.1 An Example Selection Rule

Below is a simplified version of selection rule 3.[11] It is specified in enough detail to account for the example output described earlier.

Selection Rule 3: Securities as containers for money
Parent: Rule 2
Confidence Value: 10/10
Semantic Result: $< X1 > < X2 >$ treated as physical.

$$\text{MT}: \begin{bmatrix} AP = & Purchase \\ Theme = & < X1 : Security > \\ Source = & < X2 > \end{bmatrix}$$

$$\text{OT}: \begin{bmatrix} AP = & Put-Into \\ Theme = & < X2 > \\ Goal = & < X1 > \end{bmatrix}$$

The semantic result specification in this example encodes the knowledge that the use of this selection rule and mapping schema (of describing a purchase as a put action) will cast the fillers of the *Theme* and *Source* roles as physical objects.

The confidence value in this example is very high. This means that this rule is unlikely to to yield bizarre results.

7.2 The Template Formalism

The preceding section showed an example matching schema. It had two parts, a matching template (MT) and an output template (OT). The general processing is that the input is matched against the MT.[12] Here the language in which these templates will be specified is described.[13] For the templates to be powerful enough to express the mappings the formalism must easily capture the following concepts.

[11] There is a distinction between how a selection rule is used and how it is stored. In particular, selection rules form a generalization hierarchy and some information is therefore inherited from ancestors in the hierarchy. For clarity, this rule shows some of the inherited information explicitly. Also, the mapping schema is simplified. A more complete version is described in section 7.9.

[12] Essentially during this process the variables in the MT will be bound to pieces of the input logical form. These same variables may appear in the OT where the values will be used to produce the metaphorical expression.

[13] This formalism is motivated by the functional unification formalism [Kay79].

- *And:* Square brackets represent *and*. For the MT the template will not be said to match unless all the elements in the square brackets can match with the input.

- *Or:* Curly brackets represent *or*, where a choice between options is available. Simple choices also may be represented with a shorthand of placing a || between the choices. The MT is said to match as long as one of the choices matches the input. A set of options in curly brackets may be referred to as a *choice block* or an *or block*.

- *Coindexed Variables:* We must have a link between the MT and the OT to capture the mappings. One kind of link will be variable names that are shared between the MT and OT. These *coindexed variables* will be of the form <X1>. During the matching of the MT these variables become bound. Their occurrence in the OT signifies where these values belong in the metaphor.

  These coindexed variables may be typed to specify what matching is appropriate. The MT of selection rule 3, on page 27, includes the expression: "Theme = $< X1 : Security >$." This signifies (1) that the variable X1 is to fill the role of *Theme* and (2) that the variable X1 is required to be of type *Security* (as in stocks and bonds). The variable types are used in the matching process. Thus, using rule 3, if the input included "*Theme* = tree", the matching would fail because *tree* is not a kind of security.

- *Coindexed Choices:* We also need a link between the MT and OT when it comes to the the choices that are selected (e.g., in the case where an *or* choice was available in the MT). If a particular choice is made in filling the MT, we may need to specify a choice in the OT that is dependent on the MT choice. The *coindexed choice variables* provide that link. They will be of the form <C1> and are merely index markers. These variables are not bound to any item in the input. Rather, one can think of the coindexed choice variables as instantiated to True if the associated choice is taken in the MT. It is instantiated to False otherwise. When the OT is instantiated based on the matched MT, only the coindexed choices that contain variables set to True will be taken. The complete version of selection rule 3, on page 35, contains an example of an *or block* with coindexed choice variables. In that example, if the input contains an AP of type "purchase", the output will contain a "put-into" AP.

- *Role-Filling Specifications:*

  We must be able to specify whether a role must be filled or not to allow the matching templates to be more flexible so that they only need to specify elements of the input that are actually affected by the metaphor. So far we have assumed that each role mentioned in the MT must occur in the input. Additional roles may occur in the input (which the MT does not refer to). These additional roles will be carried over to the output unaltered. The matching process is described in detail below. We need to constrain further the matching process by 1) requiring some roles in the input (this is the default), 2) forbidding some roles in the input, and 3) specifying that some roles are optional. Such specifications are accomplished through the use of the reserved word *NONE* and the use of choice blocks.

If a role in the MT has a variable of type NONE specified, that role must not occur in the input for matching to succeed. Optionality can be represented through the use of a *choice block*. This choice block must include a choice containing the optional role specified as type NONE and another choice that specifies the role as a fillable type.

7.3 The Matching Process

The matching process tests for compatibility and also binds items from the input to variables in the MT and OT. The term *matching* may be misleading because items in the input not instantiated to a variable in the MT will be maintained in the resulting structure. Matching shares this inclusive property with traditional set union.

This formalism allows for simple processing. First the input is matched to the MT. The way the input is matched to the MT is a special kind of union of the structures. We can think of each role specification as an attribute value pair. In the *Theme* part of the MT in the previous example (Theme = $< X1 : Security >$), *Theme* is the attribute and $< X1 : Security >$ is the value.

As this matching of the input to the MT is occurring, the coindexed information between the MT and the OT will be used to update the values of the coindexed variables. Also, roles and modifiers that are found in the input but not specified in the MT will be copied directly to the OT. There is good reason for this. The generative power of rule 3 should not depend one way or the other on whether the stock purchase is being made for the speaker's grandmother. In other words, the rule should work whether the input structure specifies a *Benef* or not.

7.4 General Knowledge Base

Even in a simple example we will need access to some general knowledge. To account for the example given on page 20, we must have available the knowledge that a stock is a security if the matching between rule 3 and the example input is to be successful. Rule 3 is written in terms of the more general *security*, while the given input is written in terms of the more specific *stock*. By keeping the rule as general as possible, it will work just as well with the purchase of a bond. The general knowledge must be used during the matching process. We assume a hierarchical knowledge base where concepts that may occur in an input specification are classified (possibly along several different dimensions). The knowledge base currently consists of types and tokens related with *is-a* links. For instance the knowledge base will contain the following to account for the match shown in the first example.

```
stock is-a security
bond is-a security
```

7.5 Achieving the Goals

The goal of conceptual fit is unique because every metaphor generated should have an appealing conceptual fit, even if this is not specifically requested. This is because a selection rule should only be in the system if it adheres to the conceptual roots principle. By being listed in the hierarchy, a rule is understood to be promising for the goal of conceptual fit. Also salient to the goal of conceptual fit is how the the metaphor manipulates the objects with respect to the concreteness principle.

The goal of brevity need not add any more complexity to the model. When the candidate metaphorical domains are being considered, and we know what specific role needs to be filled, we can compare the length of the specialized lexical expressions once they are identified. Although the scope of our work does not cross into the final phase of sentence generation, we could merely pass several candidates as output, and the final sentence generator could temporarily generate them so their lengths can be compared.

The goal of raising the focus level of a particular element of an expression would take advantage of the work described in section 4.1, especially the semantic focusing principle. In particular, candidate metaphorical domains must fill focusing criteria. The element to be raised must be shifted to a semantic type that is higher in the focus hierarchy than is the case in the more literal form. If many candidates are promising, we would then want to choose one that could allow the possibility of the element whose focus is to be raised to be located in subject position.

7.6 Processing the Example

Consider the input found on page 20 being matched with selection rule 3. In the first step of the processing, the input is matched to the matching template in selection rule 3. The variables are instantiated into the MT, and therefore to the coindexed variables in the OT. If the matching did not succeed (as would be the case if the input *Theme* was *tree*) processing with this selection rule would terminate.

The resulting intermediate structure would look like the following:

$$\text{OT'}: \begin{bmatrix} AP = & Put - Into \\ Theme = & Dollars(Quantity: 8000) \\ Goal = & Stock(Quantity: 100; Company: Exxon) \end{bmatrix}$$

Because the MT is successfully matched, it can potentially be used to create a metaphor. Next, the rule must be scrutinized with respect to the input goal.[14] The goal in this example is to provide a conceptual fit. By being listed in the hierarchy, a rule is understood to be promising for this goal because each selection rule must adhere to the conceptual roots principle (by definition). The information about the semantic result tells the system that selection rule 3 also adheres to the concreteness principle. Finally, the confidence value is considered. Its CV is the highest possible. Because this rule looks promising, the OT used to yield the example output.

7.7 T-M Metaphors and this Representation

In section 3 transparently-motivated metaphors are described as having four qualities (i.e., special lexical expressions, tenor anchor, universal groundings, and subtlety). These four qualities will be assured by placing stipulations on the selection rules.

T-M metaphors have specialized lexical expressions in the metaphorical domain for the action being described. This expression conveys the meaning of the action taking place. For the T-M metaphors, the verb conveys the domain (and thus accounts for the special lexical expression). Because of this, the change of verb is fundamental to generating T-M metaphors. This is accounted for through the AP. Each AP comes from a single domain (although that domain may have more general ancestors). The mapping schema will specify which domains are being accessed via the APs of the matching and output templates. The AP of the OT specifies the metaphorical domain to the surface generator. It is then the generator's job to provide an expression that conveys the meaning of the AP in a very *literal* way (in terms of the metaphorical domain). To insure that T-M metaphors are generated, each selection rule must satisfy the following stipulation. A selection rule's AP in the MT must be from a different domain than the AP in the OT.

T-M metaphors also contain a tenor anchor, which refers to an item that exists in the tenor domain and occurs in the metaphor. In this representation, the tenor anchor will be accounted for by insuring that a variable filling a role in the MT will be repeated for some role of the OT. To insure that T-M metaphors are generated, each selection rule must successfully instantiate at least one coindexed variable.

T-M metaphors are based on universal groundings that are often linked to bodily experience. The selection rules have the responsibility for maintaining this quality. When selection rules are developed, the motivation of universal groundings must be kept in mind.

The final quality of T-M metaphors is that they are subtle in the way that they do not draw

[14] Depending on the distribution of the selection rule hierarchy and the qualities of the goal, it may be more efficient to consider the goal *before* testing the match. This might be the case if the goal could disqualify a large enough portion of the selection rules. However, the selection rule hierarchy will have the property that if a rule cannot accept the input, neither will any of its descendants. Therefore, the average depth of the selection rules is crucial to determining which approach is most efficient. If the average depth of a node is minimal (i.e., there is significant breadth to the tree), it may be more efficient to check goal satisfaction before matching.

attention to themselves as non-literal. This is, perhaps, the most difficult of the qualities to assure. There is no single structure that has the responsibility for insuring that the generated metaphors are subtle. Rather, the structures accounting for the other three qualities, along with the context, combine to have the effect of subtlety.

7.8 Focus Example

It may be useful to see another example. This time, the goal will be that of manipulating focus. Below is an input specification followed by the output that the metaphor generation component will produce. The input can be glossed as: "Convey that John produced an idea while achieving the goal of placing greater focus on the idea." The output of the system could be generated as, "The idea hit John."

Input Goal: Focus(idea)

Input Message:

- **AP**: Create
  - **Theme:** Idea
  - **Agent:** John

Output Message:

- **AP**: Hit
  - **Theme:** John
  - **Agent:** Idea

This change from input message to output message is accomplished through the matching of Selection Rule 5. Below it is specified in enough detail to account for the example output.[15]

Selection Rule 5: Concept physically demanding attention

Parent:	None
Confidence Value:	7/10
Semantic Result:	$< X1 >$ treated as animate.

[15] There is room for generalization. *Create* is not the only appropriate matching AP (e.g., *have*) and *hit* is not the only output AP (e.g., *accost, find*). This probably will be encoded with an entry in the general knowledge base specifying a more general AP that the appropriate APs can share as an ancestor. However, this also could be encoded through several choices in this rule, or several similar mapping rules

$$\text{MT}: \begin{bmatrix} AP = & Create \\ Theme = & < X1 > \\ Source = & < X2 : Concept > \end{bmatrix}$$

$$\text{OT}: \begin{bmatrix} AP = & Hit \\ Agent = & < X2 > \\ Theme = & < X1 > \end{bmatrix}$$

General Knowledge base

`idea is-a concept`

First the Selection rule is successfully matched to the input. Then the goal is considered. The semantic result specification is very important when the goal is raising the focus of an item. We must know two qualities to reason about the control of focus of an item: the "normal" semantic focus value of the item and the new metaphorical focus value. The semantic result specification tells us where in the focus hierarchy the item will be if the rule is matched. If this new value is greater than the "normal" value, then metaphor can semantically raise its focus. If other rules fared equally well, the CV would be compared. The structure generated by the rule with highest CV would be output.

7.9 Hierarchical Structure

When representing knowledge it is almost always useful to to capture generalities. There is a distinction between how a selection rule is used and how it is stored. Selection rules will be stored in a generalization hierarchy and some information for a particular selection rule may be inherited from ancestors in the hierarchy. In this way, general information can be represented at its most appropriate level.

Rule 3's parent is rule 2 (on page 35). Rule 2 is entitled "Things as containers of their creation." Example input and output for rule 2 would be "A lot of his time created that relationship," and "A lot of his time went into that relationship," respectively. The matching templates grow more restrictive as we proceed down the hierarchy. This quality will allow for an efficient search for candidate selection rules.

It initially seemed reasonable to develop a "top-level" ancestor for rules 2 and 3 to capture the basic generalities under the title of "containment." However, it proved to be frustrating to specify a rule at such a general level that would be able to capture useful generalizations about its child rules and be capable of yielding metaphors itself. It became clear that the approach was combining two different kinds of information in the hierarchy–tenor domain information and metaphorical domain information.

The input, which is based on the tenor domain, is matched against the MTs of the selection rules. One reason the rules are stored hierarchically is to facilitate the search. It does this by insuring that child rules are more restrictive in matching, so that they are only appropriate to match with the input if their ancestors are. This nature can be seen in the relationship between rules 2 and 3.

A selection rule should be able to inherit information both from the tenor side and from the metaphorical side. On the tenor side, a selection may inherit information from the tenor domain via its ancestor rules. Because the MT and OT are so closely related by coindexed variables, both the MT and OT will be influenced by the ancestor selection rules. On the metaphorical side, the selection rule indicates its metaphorical domain via the AP in the OT. We need a *domain hierarchy* to capture the generalizations between domains. This captures what the failed attempt at a general rule for containment was intended to capture. The AP in the OT will identify the selection rule's metaphorical domain. This domain will be identified in the *domain hierarchy.* This domain, or one of its ancestors will indicate needed information. The *semantic result* field (which tells how the word meanings are shifted by the metaphor) is currently the only field that is inherited from the metaphorical domain.

The representation of the *domain hierarchy* will (at least at first) be very sparse. It will explicitly identify APs with their appropriate domain labels, like *containment.* The domain labels also will indicate the semantic result associated with their use in the OT. The *domain hierarchy* also provides a vocabulary of AP-types that can be referenced inside the MTs and OTs. This will be useful when specific APs are not broad enough to cover what the template is intended to represent. Below is an example of what will be found in the *domain hierarchy.*

```
Put-Into is-a Containment-AP
Take-Out-of is-a Containment-AP
Containment-AP has (Semantic Result:  (Theme:  physical; Goal:  physical))
```

The benefits of the *domain hierarchy* are not limited to the information that it stores. It also will be useful in providing guidelines for selection rule development. First, it provides the representational distinction between the tenor and metaphorical domains. This should limit problems with building the hierarchies, such as mistakenly positing a "top-level" rule for "containment." Second, it also may lend coherence to the semantics of the very flexible roles within domains. This is because the semantic shift information will partially specify the semantics of some of the roles in the various domains.

Here are the complete specifications for rules 2 and 3. Recall, rule 2 accounts for "He put a lot of himself into that relationship."[16]

[16] The *domain hierarchy* will include: `Purchase-AP is-a Possess-AP` and `Sell-AP is-a Possess-AP`.

Rule 2: Things as containers of their creation

Parent: Rule 1

Confidence Value: 4/10

Semantic Result: $< X1 > < X2 >$ treated as physical. (same as parent)

$$\text{MT}: \begin{bmatrix} AP = & < Create \parallel Maintain \parallel Possess > \\ Theme = & < X1 > \\ Instr = & < X2 > \\ Goal = & < X3 : NONE > \end{bmatrix}$$

$$\text{OT}: \begin{bmatrix} AP = & Containment - AP \\ Theme = & < X2 > \\ Goal = & < X1 > \end{bmatrix}$$

Rule 3: Securities as containers for money

Parent: Rule 2

Confidence Value: 10/10

Semantic Result: $< X1 > < X2 >$ treated as physical. (same as parent)

$$\text{MT}: \begin{bmatrix} \text{AP} = & \begin{Bmatrix} Purchase - AP; < C1 > \\ Sell - AP; < C2 > \end{Bmatrix} \\ Theme = & < X1 : Security > \\ Instr = & < X2 : Money > \\ Goal = & < X3 : NONE > \end{bmatrix}$$

$$\text{OT}: \begin{bmatrix} \text{AP} = & \begin{Bmatrix} Put - Into; < C1 > \\ Take - Out - of; < C2 > \end{Bmatrix} \\ Theme = & < X2 > \\ Goal = & < X1 > \end{bmatrix}$$

8 Application to Augmentative Communication

Many ideas in this work apply to the problems of augmentative communication, the field of developing tools to facilitate the expression of disabled people. A large population of non-speaking, motor-impaired individuals must rely on computer-based communication aids. These aids, in essence, present the user with a list of letters, words, or phrases that may be selected to compose the desired utterance. The resulting utterance may then be passed to a document preparation system. While many of these individuals desire to communicate in complete well-formed sentences, the expenditure in effort and time is often prohibitive.

The *sentence compansion system* is a project being implemented at the Applied Science and Engineering Laboratories of the AI duPont Institute and the University of Delaware. The purpose of the system is to increase the communication rate of disabled people by employing natural language processing techniques. A parallel goal of the project is to do this while placing as small a cognitive burden on the user as possible. Therefore, requiring the user to memorize a complicated encoding scheme is not an option. As input it takes only the uninflected content words of a sentence (e.g., John go store yesterday) and generates a complete well-formed sentence (e.g., "John went to the store yesterday"). It can take a disabled person selecting words from a word selection system a significant amount of time to select each word, so every word that is saved represents a valuable speed up. The name *compansion* comes from the idea that the input to the system is *comp*ressed, and that its output is an exp*ansion* of the input that yields the desired sentence. A major task of the system is to figure out what is intended by the user. It cannot rely upon conventional parsing methods, because of its telegraphic input. Instead it relies upon semantic information associated with the words [MDJP90], [JDMP91].

This work on metaphor was first motivated by an analysis of where further research on the sentence compansion project should go. I wished to determine what improvements were needed in the system's semantic component (my particular responsibility). I used the simple but effective investigation technique of studying naturally occurring examples of text. Studying natural text with a specific computer-based system in mind is a great way to make one sensitive to the prevalence of non-literal expressions. The success of the present system is marred by one major shortcoming. The sentence compansion system (the semantic component, in particular) always had associated with it the caveat that it was intended to only handle literal expressions. It seemed harmless enough to rule out the generation of poetry. However, as this work reflects, not only are non-literal uses of language prevalent, they are often difficult to distinguish. This proposed research applies directly to the goal of minimizing the cognitive load because to require the user to recognize transparently-motivated metaphors would place considerable cognitive load on the user. In other words, it would be an unreasonable burden to require the users of the sentence compansion system to consciously avoid using such expressions [JMD92].

As the name suggests, the sentence compansion system works at the sentence level. Its telegraphic input consists of words which contribute to the formation of a single sentence. Its output is that desired sentence. The work proposed here will be limited to the sentence level. This concentration is compatible with the goal of expanding the coverage of the sentence compansion system outside the purely literal realm.

It is important to note, for two reasons, that most of the knowledge in the sentence compansion system's semantic parser is about verbs. In particular, it contains knowledge about what roles of the verb can/must be filled, and what types of things those roles are most reasonably filled with. Because predictive verb knowledge is fundamentally important to the semantic parser, the usefulness of this research is enhanced by the decision to concentrate on verb phrase metaphors. The second reason is most clearly stated through an alternative solution.

Without recognizing general principles behind metaphor, the only other option to increase the coverage of the sentence compansion system would be to create extra metaphorical senses of the verbs. Such an approach would demand that the verb *put* have a special sense. The THEME role would prefer to be filled with a sum of money, and the GOAL would prefer to be filled with an investment. This might be a reasonable approach if the system was to be used within a limited domain. But since this is not the case, such an approach would clearly lead to an explosion in the number of verb senses needed. This would cause the sentence compansion system to generate many more incorrect guesses at the user's intended meaning. The user would unfortunately have to contend with these extra interpretations. Such an approach is especially foolish when it is clear that general structures for this behavior exist.

Since the users of augmentative communication systems communicate at a slow rate, it is reasonable to assume that many such individuals employ strategies that are economical with respect to the number of selections required (which relates to utterance length). As we have seen, metaphor can help accomplish the goal of being brief. It does this both by simply allowing more surface options, and by allowing large chunks of information to be conveyed economically. Therefore, given the capacity to generate metaphor, users of augmentative communication systems may express themselves in metaphor even more frequently than the general public.

Generally, a greater understanding of a relevant domain will enhance the performance of any prediction for rate-enhancement. A straight-forward example of this is statistical word prediction. In word prediction, a system may intend to make a prediction about what word is intended from the first character(s). This cannot be accomplished without a wealth of information, particularly statistics about word frequencies. Understanding metaphors more deeply from a formal, computational approach will benefit the general task of predicting such expressions. Stipulating what is reasonable and what is not reasonable in this metaphorical realm, a realm that is entered frequently and often unconsciously, can only help prediction of such expressions.

9 Evaluation Criteria

Here I consider how to demonstrate that the proposed metaphor generation system will generate appropriate metaphors that achieve the requested goal. There are two parts to evaluating the success (desirability or appropriateness) of a generated metaphor. The most important part is that the basic meaning can be easily recovered. The other part is the achievement of the goal. The basic meaning will be considered first.

The most specific rules (with high CVs) should yield metaphors that are easily understood (ones with low CVs will be less reliable). However, due to the murky nature of metaphor, and because this research is the first attempt at generating this class of metaphors, rather than a new attempt at a previously studied problem, some failure must be anticipated. When the system fails by generating a metaphor that is clearly not conveying the intended meaning, I must be able to explain why this failure happens and what knowledge the system would need to avoid such undesirable output.

To judge the acceptability of the generated metaphors, we must first generate English sentences that would result from the metaphor generation component's conceptual output. The best way to judge the acceptability of these metaphors is to have native English speakers interpret the meanings of these sentences. If the judges can determine the meaning of these sentences success can be claimed.

We can consider understanding such metaphors as having two distinguishable parts: identifying the tenor domain, and identifying the meaning in the tenor domain. We must keep in mind, however, that the judges must not expect all expressions to sound familiar. The question is, can they be understood properly, and not be misleading or incomprehensible.

Lets now turn to the other part of evaluating the success of these metaphors, testing the goal was achieved. Although testing the goal of brevity may be quite simple, testing the more complex goals of conceptual fit and focus control present serious problems.

The question of testing the success of focus control brings to mind the distinction that this work is aimed at providing metaphor as a useful *option* to a text planner. The question of when it is appropriate to use a metaphor is a large open question beyond the scope of this work. Similarly, it is an open question how focus claims can be empirically tested. Hudson-D'Zmura's Ph.D. thesis [HD89] was dedicated to testing focus claims experimentally. Several problems confounded her work. Tense, the semantics of words (e.g., causality) and pragmatics can all effect the experimental measurements that have been used to measure focus. Her work was also confounding by the weakness of the focusing theory (centering). This theory can only define the focus of simple sentences.

There is currently not an accepted way to test focus. The most fundamental barrier to empirically testing focus is that focus is defined (and then only prototypically) in terms of text structure–not psychological states. Focus claims can only be tested indirectly.

Consider the simple method of testing focus claims in which we measure the time it takes for subjects to resolve pronouns. Many confounding variables could be controlled for when testing the focus changes between literal statements and their metaphorical counterparts. Tense could be held constant. With some limitations, surface order and some pragmatic concerns also could be held constant. (e.g., If the context was "John hit Bob" the *He* in "He got angry" would be preferred to be Bob, not John.) However, by the nature of these metaphors, it is unclear how to eliminate the interference of other variables. The degree of formality probably interacts profoundly with the choice to use a metaphor. Also, the semantics of the words relevant to focus change cannot be held constant. This is because the manipulation of semantics is the primary method for manipulating focus via metaphor. Such a study will not work. This attempt to measure focus control via anaphora resolution would be either useless or test something other than what we intended to test.

The best practical way to test whether the metaphor is making something higher in focus may be to isolate the semantics of the words changed by the metaphor. However, this leads to a discussion that has already been considered. The system will employ the *semantic focusing principle* (that we achieve greater focus by moving items up the focus hierarchy) when attempting to manipulating focus. We have already discussed evidence, in section 4.1, for the validity of the focusing claims related to the semantic focusing principle with respect to tests of semantics in isolation. Therefore, as long as the metaphor generation component abides by the semantic focusing principle, we can be confident that the examples will hold up to the best tests of focus available (i.e., the tests of isolated semantics).

The goal of conceptual fit is also very difficult to test experimentally. Testing its success–that comprehension is aided by the metaphor–will be as difficult as testing focus claims. We can appeal to the evidence and motivations behind the concreteness and conceptual roots principles. The system is designed to to adhere to those principles. Besides, if the conceptual fit goal fails, the failure probably will show up in the simple metaphor comprehension test.

To summarize, it will be reasonable to test how native English speakers can comprehend examples generated by the system. This test may have two parts, identifying the tenor domain, and identifying the meaning in the tenor domain. However, testing the subtle goals would be very problematic.

10 Current Direction: AP Modifiers

This section describes preliminary work on handling modifiers of the APs. The structures outlined here are not sufficiently developed to be included in the general explanation of the system. A study of modifiers of the APs and how they are mapped between domains are needed to account for metaphors such as (5) and (6).

(5) "The price soared."

(6) "The hours crept."

An important contributor to the meaning of (5) is the manner in which the price (metaphorically) moved. It will be useful to distinguish modifiers of the APs and not require distinct APs for each different manner in which the AP may be performed (e.g., one AP for moving quickly and another for moving slowly). Distinguishing this information will allow for a smaller set of APs. More importantly, as these examples indicate, the modifiers of APs may be fundamental to the meaning. If such modifiers were hidden within the semantics of the various APs, there would be no general way for the mapping structures to manipulate the AP modifiers.

The input to the metaphor generation system for example (5) would be (**AP**: Scalar-Change (*kind:* increase; *degree:* great) *Theme:* Price). The output would be of the form: (**AP**: Move (*kind:* up; *degree:* great) *Theme:* Price). The APs in (5) contains two modifiers. The value of *degree* remains constant across domains. The value of *kind* does not. To be meaningful in the metaphorical domain *kind* must be changed from increase to up. This shows that we would need to have available the information that encodes Lakoff and Johnson's [LJ80] famous observation "More is up, Less is down." [17]

Example (6) demands one AP-modifier mapping, but it is a complex mapping. The input to the metaphor generation system for example (6) would be (**AP**: Time-Pass (*Quality:* low)). The output would be of the form: (**AP**: Move (*rate:* slow) *Theme:* Time). Information that conveys such a mapping of AP modifiers would help account for how (6) could be generated. Although the modifier of an AP can be realized independently in an expression as an adverb, modifiers are concepts independent of surface structure. Example (6) uses a single word (i.e., "crept") that contains both the AP and the modifier.

10.1 Relevant Issues Beyond The Scope of This Work

Two issues outside the scope of this work should be acknowledged. In example (6), the word *crept* carries the notions of *move* and *slowly.* Unfortunately, it also can be interpreted as carrying some troubling information that may need to be ignored by the mapping. For example, we want the system to ignore the inference that the creeper (e.g., hour) has arms and legs. In a literal sense *crept* and *flew* are not the same action at all, but when used in a metaphorical way, they are the same general AP (i.e., *move*). The system will need to have a representation of the information that is to be conveyed, but it is beyond the scope of this work to reason about all possible inferences that may, or may not be invited by an expression.

The second issue is about modifiers of objects. Not only can APs have modifiers, so can the objects that fill the roles. For example, in the (non-TM) expression "That is a big problem," the *significance* of the problem is mapped to *size* in the metaphorical domain of physical space. Because T-M metaphors are based on the verb, it is important to capture the modifiers of the verbs. Although the approach to handing modifiers of APs probably will be of use to handling modifiers of objects as well, I intend to concentrate only on the AP issue.

[17]The mapping from kind increase to kind up may come from more general structures such as the concept of stacking. In the (perhaps cognitively fundamental) domain of piling or stacking, "an **increase** in stuff means the pile goes **up**."

10.2 Representing AP Modifiers

A key issue in the representation of AP modifiers is the modifier types themselves. The modifier types *kind, degree, rate,* and *manner* were used as general types of AP modifiers for the previous examples. The nature of these types is an open question. I will probably need to come up with a small vocabulary for these modifiers.

It is unclear whether the modifier-type "kind" is specific enough. Perhaps an appeal to vectors may be useful. Perhaps APs in most domains can be seen as having a magnitude (such as quantity) and a direction (such as up or increase). The choice of modifier types is not arbitrary. Depending on the choice of modifier types, we may or may not require a change in modifier types and/or values between domains.

It is difficult to think of some modifiers of APs in a truly literal way. For example, to think of progress in terms of something more literal than motion is quite difficult. To describe forward or backward progress as positive or negative is an attempt to strip some of the metaphor away, but it is not completely gone. Perhaps negative progress should be seen as another AP altogether. This would remove the problem of the representation being tainted by including a metaphor, but it would remove a generalization. Currently, it seems that the meaning of the AP progress should be assumed to be what we mean by forward progress. If we are to mean the opposite of that, we will use the modifier *Antithesis.* This will allow us the power of having an AP for the concept of "anti-progress" while maintaining a parsimonious representation.

A selection rule, or set of selection rules will be needed to capture the mapping *Progress is understood in terms of motion.* This will map APs about progress to APs about physical motion (e.g., "My work is finally getting somewhere"). When APs about progress (or actions in general) are mapped to motion many observations about mapping of modifiers of the APs are relevant [LESG91]. The following observations may be useful to such selection rules.[18]

- `Forward Progress is Forward Motion:` *The project is going ahead as planned.*
- `Negative Progress is Backward Motion:` *She found herself back at the beginning of the writing process again.*
- `Careful Action is Careful Motion:` *When I'm around her I feel like I am walking on eggshells.*
- `Easy Action is Easy Motion:` *It's all downhill for the rest of the course.*
- `Speed of Progress is speed of Motion:` *He flew through his work.*

[18] To be able to take advantage of many of these structures, we may require an oracle that can tell us how to express a modifier in the metaphorical domain E.g., This would reveal that careful motion can be expressed as "walking on eggshells."

- `Lack of Change Resulting from Action is Lack of Motion:` *This is getting us nowhere.*

It appears that the modifiers of APs should not explicitly be part of the selection rules.[19] The generalities between the two seem to cut across different dimensions. Rather, we will need a set of "charob" modifier routines that are triggered by the appropriate tenor and metaphorical domains.

If one of these "charobs" is applied successfully, it may be appropriate to increase the confidence value of the expression above that is specified by the selection rule. This would be the one way of providing the much-needed additional information that will allow the very general selection rules to yield metaphors. Without such additional information, the only use for selection rules with very low confidence values is as a device to help store the rule information parsimoniously.

11 Plan

I intend to develop a metaphor generation component as described in this document. There are several major tasks that must be completed to achieve this goal.

1. *Usefulness of Metaphor Generation*
   The question of when it is desirable to employ a metaphorical construction is fundamental to metaphor generation research. The answer to this has been started on. The goals described in section 4 show several specific goals that metaphor generation can be applied to. Still, it would no doubt be beneficial to look over more naturally occurring texts at several points of progress along the way. It is possible that more goals may be found and other insights reached through further text analysis.

2. *Specify the General Knowledge Representation*
   The knowledge representation at all levels must be defined more formally. I have determine the form of the input and output. The form of selection rules is well-specified, but the scope of what they should contain is not completely finished. I still need to specify what kind of additional knowledge this system will need to access (e.g., domain knowledge).

3. *Specify the Mapping Knowledge*
   The specific question of the structure of the mapping knowledge is a central issue in this work. This large problem has several parts.

   - *Add fields to the Selection Rules*
     As work proceeds additional fields may be identified that will be useful to the system.
   - *Develop a Library of Selection Rules*
     I must develop a library of selection rules for two main reasons. To be able to develop

[19]The first of that list is an exception. It contains only the default value of progress (i.e., without Antithesis)

an implementation that demonstrates the power of my approach, it must have available a set of selection rules that cover a variety of tenor and metaphorical domains. The second reason is related to the hierarchical structure.

- *Generalizing Knowledge*

  If the hierarchical structure is going to "emerge" from specific and generative selection rules, a significant number of selection rules will be required. The work on adding selection rules should continue until the characteristics of the hierarchies stabilize.

  Related to generalizing the mapping information is the promise of using general knowledge and reasoning from metaphorical domains. I may be prepared to begin specifying what kind of questions a metaphor generation system would need to ask of a domain expert.

4. *Implementation*

   After the structures have been sufficiently developed, I will put together an implementation to demonstrate the validity of the theory. The forms of the input that the system will take and its output are described in section 6.1.

5. *Relation to Augmentative Communication*

   After the implementation is well under way, it will be clear how this investigation can be used to achieve computational results. It will be possible to specify how the developed ideas and structures can, in a direct way, apply to augmentative communication. In particular, I will describe how these ideas can be incorporated with the sentence compansion project, which was described in section 8.

12 Conclusion

This paper has introduced transparently-motivated metaphor. It has shown that generation of T-M metaphors will contribute to the disciplines of both natural language processing and augmentative communication. It has shown that T-M metaphors are prevalent and that it is useful to generate such metaphors. This has described many textual goals that such metaphor can achieve and discussed principles that can explain many uses of metaphor found in the text analysis. This analysis and the resulting principles have led to developments toward a methodology for generating transparently-motivated metaphors from general principles. This accounts for the conceptual structures behind such expressions, and accounts for the effects of such manipulation.

13 Appendix A

The set of examples that follow will hopefully clarify what transparently-motivated metaphors are and the goals that can be achieved through them.

Metaphor can be used to achieve several textual goals of a computational generation system. It is not, however, claimed that people consciously use metaphor for such goals. Consider the goal of managing focus to make text more coherent. A sophisticated natural language generation system must explicitly make focus choices to partially satisfy its purpose of making its text coherent to a human audience. It is not necessarily appropriate to speak of people as using textual devices to be coherent, since being human greatly contributes to being understandable to other humans. When a person says or writes expressions, something like "the focus" is already somewhere in the speaker's head. It is more accurate to say that the focus analysis of a person's utterance *reveals* the way the speaker conceptualized the ideas that were expressed, rather than that particular choices were made.

When examining naturally occurring examples, it is not necessarily appropriate to speak in the terms that apply to computational generation systems. Included in the examples below are goals that the examples could possibly apply to. This is one reason multiple goals may be specified for a single example. This also highlights the fact that many goals are compatible with each other. A good case in point is conceptual fit. Conceptual fit applies to most T-M metaphors, to some degree.

1. "AI *is no stranger to* object-oriented paradigms." [Ell91]

   - `Purpose`
     - Focus: Raise AI to more animate
     - Brevity
   - `More Literal`: "AI and object-oriented paradigms have long been associated with each other", "incorporated together"
   - `Metaphorical Domain`: People

2. "COBOL and object-oriented approaches can be *mated* in several ways." [Ell91]

   - `Purpose`
     - Focus: Raise the approaches to more animate
   - `More Literal`: "used together", "incorporated together"
   - `Metaphorical Domain`: People

3. "The field helped *spawn* the object-oriented movement" [Ell91]

- `Purpose`
    - Focus: Raise AI to more animate. Similar to 1.
- `More Literal`: "start", "begin"
- `Metaphorical Domain`: Reproducing Animates

4. "Lisp has *adopted* various object-oriented advances" [Ell91]
    - `Purpose`
        - Focus: Raise AI to more animate. Similar to 1.
        - Brevity
    - `More Literal`: "have been incorporated into Lisp"
    - `Metaphorical Domain`: Human
    - `Note`: This could be viewed as a metonymy where Lisp signifies the developers of Lisp.

5. "Considered an object class *in* the object-oriented approach" [Ell91]
    - `Purpose`
        - Conceptual Fit: The concreteness and conceptual roots principles both apply.
        - Brevity
    - `More Literal`: "with respect to the object-oriented approach"
    - `Metaphorical Domain`: Containment

6. "C++ and object-oriented versions of Pascal must be credited for *keeping the object-oriented programming alive*" [Ell91]
    - `Purpose`
        - Conceptual Fit: A selection rule to fit this concept would convey that the usefulness of a system is its life; when it is no longer used it dies.
        - Focus: Object-oriented programming is raised to an animate level.
        - Brevity
        - Perspective The very use of life and death terms may cast a strong perspective on the matter.
    - `More Literal`: "keeping object-oriented programming pragmatically viable"
    - `Metaphorical Domain`: Animate
    - `Note`: It is very tough to capture this meaning more literally.

7. " . . . *wreck* lives . . . " and " . . . *control* lives . . . " [Kin88]

   - `Purpose`
     - Conceptual Fit: A life is not itself a physical thing that can be literally wrecked (controlled, sure). Yet, by adhering to the concreteness principle, and perhaps the conceptual roots principle, these are a very clear expressions.
     - Focus: This moves abstract up to physical.
   - `More Literal:` " . . . reduce the quality of the life of . . . "
   - `Metaphorical Domain:` Physical (perhaps moving vehicle, or physical structure). When we start thinking about wrecking a life the way we can a motor vehicle, then it is easier to view the expression "control lives" as metaphorical.

8. "The next COBOL standard should *emerge* in the mid 1990s." [Ell91]

   - `Purpose`
     - Conceptual Fit: The concreteness and conceptual roots principles both apply.
     - Focus: We lose the notion of the subject being created by a group of people, instead the focus is upon it emerging.
     - Brevity
   - `More Literal:` " . . . become available . . . "
   - `Metaphorical Domain:` Containment / Visibility

9. " Anything that can be *crammed into* an object oriented framework– including AI." [Ell91]

   - `Purpose`
     - Conceptual Fit: The concreteness and conceptual roots principles both apply.
     - Brevity
     - Perspective: This allows for use of the word *cram*, which may have a useful connotation.
   - `More Literal:` " . . . make use of . . . "
   - `Metaphorical Domain:` Containment

10. " . . . incorporating object-oriented features directly *into* COBOL." [Ell91]

    - `Purpose`
      - Conceptual Fit: The concreteness and conceptual roots principles both apply.

- Lexical hole: It is tough to come up with a more "literal" way of expressing this. That is why it is tempting to call this use of *into* a specific word sense. Note, however that this approach would account for this phenomena in a more general manner.

- `More Literal:` " . . . directly with COBOL . . . (?)"
- `Metaphorical Domain:` Containment

11. " . . . the ANSI group has *arrived* at some . . . " [Ell91]

    - `Purpose`
      - Conceptual Fit: The concreteness and conceptual roots principles both apply.
      - Focus: Raises what has been arrived at as more physical.
    - `More Literal:` " . . . (nearly) completed some . . . "
    - `Metaphorical Domain:` Motion, Journey

12. " . . . have *reached* a tentative agreement." [Ell91]

    - `Purpose`
      - Conceptual Fit: The concreteness and conceptual roots principles both apply.
      - Focus: Raises the agreement to more physical.
    - `More Literal:` "negotiated", "developed"
    - `Metaphorical Domain:` Motion, Journey or Grasping

13. "The group still has a way to go . . . " [Ell91]

    - `Purpose`
      - Conceptual Fit: The concreteness and conceptual roots principles both apply.
      - Focus: Raises the work ahead to more physical.
    - `More Literal:` "has remaining work to complete", "has not yet finished"
    - `Metaphorical Domain:` Motion, Journey

14. " . . . at the *heart* of the debate . . . " [Kin88]

    - `Purpose`
      - Conceptual Fit: The concreteness and conceptual roots principles both apply.
      - Focus: Raises the debate to more physical.
    - `More Literal:` " . . . fundamental to the debate . . . "
    - `Metaphorical Domain:` The body.

15. " . . . *channel* desires . . . " [Kin88]

- `Purpose`
  - Conceptual Fit: A desire is not itself a physical thing that can be literally channeled. Yet, by adhering to the concreteness principle, and perhaps the conceptual roots principle, this is a very clear expression.
  - Focus: This moves abstract up to physical.
  - Lexical Hole: "redirect" is also physical.
- `More Literal:` "?"
- `Metaphorical Domain:` Physical movement

16. "A baseball player can *enter* a game "[20]

- `Purpose`
  - Conceptual Fit: Although the field on which the game is played can be entered, this would be considered an inappropriate intrusion if done by a spectator or unauthorized player. entering the game means something more.
  - Brevity
- `More Literal:` "begin to participate"
- `Metaphorical Domain:` Containment

17. "By getting on base in the last inning, with his team behind, a baseball player can keep the game *alive.*"

- `Purpose`
  - Perspective: The theme of life and death has a strong impact. Here, the finality of the end of the game and the importance of the game are heightened.
  - Conceptual Fit: The process of the game (or the team's effort) is related to the (easily identified with) bodily notion of life and death.
  - Brevity
- `More Literal:` " . . . keep their chances for victory possible.", " . . . not end the game."
- `Metaphorical Domain:` Animates

18. "How can Buchanan keep his campaign *alive?*"

[20] If there is no particular reference cited, it is a spoken quote, or paraphrase which maintains the integrity of the metaphor, for which I have no specific reference.

- `Purpose`
  - Perspective: The theme of life and death has a strong impact.
  - Conceptual Fit The process of Buchanan's political campaign is related to the (easily identified with) bodily notion of life and death.
  - Brevity
- `More Literal`: " . . . keep chances for victory possible."
- `Metaphorical Domain`: Animates

19. "Buchanan must get up from the canvas (about his loss in an early primary)."

- `Purpose`
  - Focus: It is difficult to focus upon the *state* of Buchanan's campaign without a metaphor. Giving that state a physical place (on the canvas) puts the state in very high focus.
  - Perspective: Inviting the theme of a physical boxing match gives it more of the feel of a battle.
  - Conceptual Fit: Buchanan's abstract political campaign is likened to a boxer in a physical battle.
  - Brevity
- `More Literal`: " . . . must start making progress again."
- `Metaphorical Domain`: Physical battle, Boxing

20. "Buchanan has drawn blood (from Bush)"

- `Purpose`
  - Focus: The use of the word blood leaves a physical element that can be pictured to represent the "damage."
  - Perspective: The theme of a personal physical battle, complete with blood, is invited.
  - Conceptual Fit: The respective (abstract) political campaigns are likened to (physical) fighters.
  - Brevity
- `More Literal`: " . . . has lowered Bush's chances of victory."
- `Metaphorical Domain`: Physical battle, Boxing

21. "The day is slipping away"

- `Purpose`

- Focus: Makes time an object, higher in the focus hierarchy than abstract concepts.
- Perspective: Implies that we can control time. We cannot stop time, but we can get a better grip upon slipping objects.
- Conceptual Fit: The difficult abstract concept of time is simplified into a tangible object.
- Brevity

- `More Literal:` " (?) we are not being productive"
- `Metaphorical Domain:` Physical objects, Time a as resource

22. "He cannot grasp this idea"

- `Purpose`
  - Focus: Makes the idea an object, and therefore higher in the focus hierarchy.
  - Conceptual Fit: Same as focus.
- `More Literal:` "understand this idea", "comprehend this idea"
- `Metaphorical Domain:` Physical objects

23. "...move on to another topic"

- `Purpose`
  - Perspective: It may imply that progress is being made by moving.
  - Focus: Makes the topics places which can visited, and therefore higher in the focus hierarchy than abstract topics.
  - Conceptual Fit: Same as focus.
- `More Literal:` "begin discussing another topic"
- `Metaphorical Domain:` Physical objects

24. "We have slapped together another show."(David Letterman)

- `Purpose`
  - Perspective: Humor. Slapped together brings the connotation that it was created in a sloppy manner.
  - Focus: Makes the *Late Night with David Letterman* show (which has both physical and abstract components) completely physical.
  - Conceptual Fit: Same as focus.
- `More Literal:` " . . . created . . . "
- `Metaphorical Domain:` Physical objects

References

[And90] Richard C. Anderson. Inferences about word meanings. In Arthur C. Graesser and Gordon H. Bower, editors, *Inferences and Text Comprehension*, pages 1–16. Academic Press, New York, 1990.

[BE68] J. G. Borkowski and H. C. Eisner. Meaningfulness and abstractness in short-term memory. *Journal of Experimental Psychology*, 76:57–61, 1968.

[BFP83] S. E. Brennan, M.W. Friedman, and C. J. Pollard. A centering approach to pronouns. In *Proceedings of the 25th Annual Meeting*, pages 155–162, Stanford CA, June 1983. Association for Computational Linguistics.

[Bla62] Max Black. *Models and Metaphors.* Cornell University Press, Ithica, NY, 1962.

[Bla79] Max Black. More about metaphor. In A. Ortony, editor, *Metaphor and Thought*, pages 19–43. Cambridge University Press, Cambridge, 1979.

[Car82] Jaimie G. Carbonell. Metaphor: An inescapable phenomenon in natural language comprehension. In W. Lenhart and U. Ringle, editors, *Strategies for Natural Language Processing*, chapter 15, pages 415–434. Lawrence Erlbaum Associates, Hillsdale, NJ, 1982.

[Car87] David Carter. *Interpreting Anaphors in Natural Language Texts.* John Wiley and Sons, New York, NY, 1987.

[Cor91] Eileen Cornell Way. *Knowledge Representation and Metaphor.* Kluwer Academic Publishers, Boston, 1991.

[DB66] William F. Dukes and Jarvis Bastian. Recall of abstract and concrete words equated on meaningfulness. *Journal of Verbal Learning and Verbal Behavior*, 5:455–458, 1966.

[Dow91] David Dowty. Thematic proto-roles and argument selection. *Language*, 67(3):547–619, 1991.

[Ell91] Lance B. Elliot. The bandwagon blues. *AI Expert*, 6(5):11–13, May 1991.

[Fas91] Dan Fass. met*: A method for discriminating metonymy and metaphor by computer. *clj*, 17(1):49–90, March 1991.

[Fra92] William Frawley. *Linguistic Semantics.* Lawrence Erlbaum Associates, Hillsdale, NJ, 1992.

[FW83] Dan Fass and Yorick Wilks. Preference semantics, ill-formedness and metaphor. *American Journal of Computational Linguistics*, 9(3-4):178–187, July-Dec 1983.

[Gen83] D. Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7(2):155–170, 1983.

[Gib84] Raymond W. Gibbs, Jr. Literal meaning and psychological theory. *Cognitive Science*, 8(3):275–304, 1984.

[Gib89] Raymond W. Gibbs, Jr. Understanding and literal meaning. *Cognitive Science*, 13:243–251, 1989.

[GJW83a] B. Grosz, A.K. Joshi, and S. Weinstein. Providing a unified account of definite noun phrases in discourse. In *Proceedings of the 21st Annual Meeting*, pages 44–50, Cambridge, Mass, June 1983. Association for Computational Linguistics.

[GJW83b] Barbara J. Grosz, Aravind K. Joshi, and Scott Weinstein. Providing a unified account of definite noun phrases in discourse. In *Proceedings of the 21st Annual Meeting*, pages 44–50, Cambridge, Mass, June 1983. Association for Computational Linguistics.

[Gor61] Aloysia M. Gorman. Recognition memory for nouns as a function of abstractness and frequency. *Journal of Experimental Psychology*, 61:23–29, 1961.

[HD89] Susan Berenice Hudson-D'Zmura. *The structure of discourse and anaphor resolution: the discourse center and the roles of nouns and pronouns*. PhD thesis, University of Rochester, 1989.

[Jac85] Paul Schafran Jacobs. *A Knowledge-Based Approach to Language Production*. PhD thesis, University of California at Berkeley, 1985. Computer Science Division Report.

[JDMP91] Mark Jones, Pat Demasco, Kathleen McCoy, and Christopher Pennington. Knowledge representation considerations for a domain independent semantic parser. In *Proceedings of the 14th Annual Conference*, pages 109–111, Kansas City, MO, 1991. RESNA.

[JM92] Mark A. Jones and Kathleen F. McCoy. Transparently-motivated metaphor generation. In R. Dale, E. Hovy, D. Rosner, and O. Stock, editors, *Aspects of Automated Natural Language Generation: The 6th International Workshop on Natural Language Generation Proceedings*, pages 231–246, Trento, Italy, 1992. Springer-Verlag.

[JMD92] Mark Jones, Kathleen McCoy, and Patrick Demasco. The problem of metahpor in intelligent aac systems. In *Proceedings of RESNA International '92 Conference*, pages 363–365, Toronto, Canada, 1992. RESNA.

[Joh87] Mark Johnson. *The Body in the Mind: The Bodily Basis of Reason and Imagination.* University of Chicago press, Chicago, 1987.

[Jon92] Mark A. Jones. Generating a specific class of metaphors. In *Proceedings of the 30th Annual Meeting of the ACL*, pages 321–323, Newark, DE, 1992. ACL.

[Kay79] Martin Kay. Functional grammar. In *Proceedings of the 5th Annual Meeting.* Berkeley Linguistics Society, 1979.

[Kee85] Edward L. Keenan. Passive in the world's languages. In Tim Shopen, editor, *Language Typology and Syntactic Description*, volume 1: Clause Structure, pages 243–281. Cambridge University Press, Cambridge, 1985.

[Kin88] Mike Kinsley. Glass houses and getting stoned. *Time Magazine*, page 92, June 6 1988.

[KPG86] R. Kittredge, A. Polguere, and E. Goldberg. Synthesizing weather forecasts from formatted data. In *Proceedings of Coling86*, pages 563–565. Coling, 1986.

[Kuk83] Karen Kukich. Design of a knowledge-based report generator. In *Proceedings of the 21st Annual Meeting of the ACL*, pages 145–150, Cambridge, Ma, June 1983. ACL.

[Lak87] George Lakoff. *Women, Fire and Dangerous Things: What Categories Reveal About the Mind.* University of Chicago Press, Chicago, 1987.

[Lak93] George Lakoff. The contemporary theory of metaphor. In A. Ortony, editor, *Metaphor and Thought (Second Edition).* Cambridge University Press, Cambridge, 1993. in press.

[LESG91] George Lakoff, Jane Espenson, Alan Schwartz, and Adele Goldberg. *Master Metaphor List.* University of California at Berkely, Berkely CA, 1991.

[LJ80] George Lakoff and Mark Johnson. *Metaphors we live by.* University of Chicago Press, Chicago, 1980.

[LKB92] Steven L. Lytinen, Jeffrey D. Kirtner, and Robert R. Burridge. Literal meaning and the comprehension of metaphors. In *Proceedings of the 1986 National Conference on Artificial Intelligence*, San Jose, CA, 1992. Coling. In Press.

[Mar87a] James H. Martin. The acquisition of polysemy. In *Proceedings of the 4th International Workshop on Machine Learning*, pages 198–204. 1987.

[Mar87b] James H. Martin. Understanding new metaphors. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pages 137–139, Milan, Italy, 1987. IJCAI-87.

[Mar88] James H. Martin. Representing regularities in the metaphoric lexicon. In *Proceedings of Coling88*, pages 396–401. Coling, 1988.

[Mar90] James H. Martin. *A Computational Model of Metaphor Interpretation.* Academic Press, 1990.

[McC89] Kathleen F. McCoy. Generating context sensitive responses to object-related misconceptions. *Artificial Intelligence*, 41:157–195, 1989.

[MDJP90] Kathleen McCoy, Pat Demasco, Mark Jones, and Chistopher Pennington. Applying natural language processing techniques to augmentative communication systems. In *Proceedings of the 13th International Conference on Computational Linguistics*, pages 413–415, Helsinki, Finland, August 1990. COLING.

[NL87] Peter Norvig and George Lakoff. Taking: A study in lexical network theory. In *Proceedings of the 13th Annual Meeting*, pages 195–206. Berkeley Linguistics Society, 1987.

[Pai67] Allan Paivio. Paired-associate learning and free recall of nouns as a function of concreteness, specificity, imagery, and meaningfulness. *Psychological Reports*, 20:239–245, 1967.

[Pai79] Allan Paivio. Psychological processes in comprehension. In A. Ortony, editor, *Metaphor and Thought*, pages 150–171. Cambridge University Press, Cambridge, 1979.

[PYM68] Allan Paivio, J. C. Yuille, and S. Madigan. Concreteness, imagery, and meaningfulness values for 925 nouns. *Journal of Experimental Psychology*, 76(1), 1968. Monograph Supplement part 2.

[Ric36] Ivor A. Richards. *The Philosophy of Rhetoric.* Oxford University Press, 1936.

[Sid79] C. L. Sidner. *Towards a Computational Theory of Definite Anaphora Comprehension in English Discourse.* PhD thesis, MIT, June 1979.

[Sin90] Murray Singer. *Psychology of Language.* Lawrence Erlbaum Associates, Hillsdale, NJ, 1990.

[Sto29] Stuart M. Stoke. Memory for onomatopes. *Journal of Genetic Psychology*, 36:594–596, 1929.

[Sur92] Linda Z. Suri. Correcting illegal NP omissions using local focus. In *Proceedings of the 30th Annual Meeting (Student Session)*, pages 273–275. Association for Computational Linguistics, June 1992.

[TS82] Roger Tourangeau and Robert J. Sternberg. Understanding and appreciating metaphors. *Cognition*, 11:203–244, 1982.

[Wei85] E. Judith Weiner. Solving the containment problem for figurative language. *International Journal of Man-Machine Studies*, 23:527–537, 1985.

[Wil86] Robert Wilensky. Some problems and proposals for knowledge representation. Technical Report UCB/CSD 86/294, University of California, Berkeley, 1986.