met*: A Method for Discriminating Metonymy and Metaphor by Computer

Dan Fass*
Simon Fraser University

The met method distinguishes selected examples of metonymy from metaphor and from literalness and anomaly in short English sentences. In the met* method, literalness is distinguished because it satisfies contextual constraints that the nonliteral others all violate. Metonymy is discriminated from metaphor and anomaly in a way that [1] supports Lakoff and Johnson's (1980) view that in metonymy one entity stands for another whereas in metaphor one entity is viewed as another, [2] permits chains of metonymies (Reddy 1979), and [3] allows metonymies to co-occur with instances of either literalness, metaphor, or anomaly. Metaphor is distinguished from anomaly because the former contains a relevant analogy, unlike the latter. The met* method is part of Collative Semantics, a semantics for natural language processing, and has been implemented in a computer program called meta5. Some examples of meta5's analysis of metaphor and metonymy are given. The met* method is compared with approaches from artificial intelligence, linguistics, philosophy, and psychology.*

1. Introduction

Metaphor and metonymy are kinds of figurative language or *tropes*. Other tropes include simile, irony, understatement (litotes), and overstatement (hyperbole).

Example 1
"My car drinks gasoline" (Wilks 1978, p. 199).

Example 2
"The ham sandwich is waiting for his check" (Lakoff and Johnson 1980, p. 35).

Sentences (1) and (2) contain examples of metaphor and metonymy respectively. Neither sentence is literally true: cars do not literally drink nor do ham sandwiches literally wait. Notice, though, that the two sentences are interpreted differently. "My car" in (1) is commonly understood as resembling an animate drinker while in (2) "the ham sandwich" is generally interpreted as referring to the person who ordered the ham sandwich.

Most of the considerable literature on metaphor and the smaller one on metonymy (see Van Noppen, De Knop and Jongen 1985; Shibles 1971) is from philosophy, linguistics, and psychology. On the whole, the two phenomena remain vague, poorly defined notions in that literature. In artificial intelligence (AI), detailed treatments of either metaphor or metonymy are relatively scarce. Moreover, most of those treatments are paper implementations that have not been coded up and run on a computer.

* Centre for Systems Science, Simon Fraser University, Burnaby, British Columbia, Canada V5A 1S6

The met* (pronounced "met star") method provides a means for recognizing selected examples of metonymy and metaphor, and also anomaly and literalness, in short English sentences.[1] The method is part of Collative Semantics (hereafter CS), which is a semantics for natural language processing. CS, and hence the met* method, has been implemented in a program called meta5 (so called because it does more than metaphor). The meta5 program is, as far as I know, the first system to recognize examples of metaphor and metonymy. To my knowledge, there is only one other working program that might be said to recognize instances of metaphor (Martin 1988; 1990) and two systems that appear to recognize cases of metonymy, TEAM (Grosz et al. 1987) and TACITUS (Hobbs and Martin 1987).

The rest of the paper is organized as follows. Section 2 surveys general issues and approaches in metaphor and metonymy, notably the distinctive characteristics of metaphor and metonymy, the relationship between metaphor and metonymy, and the relationship between literalness and nonliteralness. Section 3 presents the met* method, concentrating on the basic topology of the met* method algorithm. Section 4 shows details of representations and processes used in CS. Section 5 gives examples of the meta5 program analyzing simple metaphors and metonymies. Descriptions get progressively more detailed from Section 2 through to Section 5. Sections 6 and 7 describe some extensions to metaphor interpretation in CS and compare the met* method against other approaches to metaphor and metonymy, especially computational ones. A glossary of key terms is provided at the very end of the paper.

2. Survey of Metonymy and Metaphor Research

Metonymy and metaphor are so poorly understood that widely divergent views exist about them and their relationship to each other. This section reviews research on metaphor (2.1), metonymy (2.2), the relationship between them (2.3), and the more general relationship between literalness and nonliteralness (2.4).

2.1 Metaphor

Four views of metaphor are critically discussed: the comparison view, the interactive view, the selection restriction violation view, and the conventional metaphor view. Computational examples of each kind are included by Gentner, Indurkhya, Hobbs, Wilks, and Martin. Space does not permit discussion of other AI work on metaphor by, e.g., Russell (1976) and Weiner (1984; 1985).

2.1.1 The Comparison View. According to the comparison view

> a metaphor is a comparison in which one term (the *tenor* or subject of the comparison) is asserted to bear a partial resemblance (the *ground* of the comparison) to something else (the *vehicle*), the resemblance being insufficient to sustain a literal comparison. As with any comparison, there is always some residual dissimilarity (the *tension*) between the terms involved in the comparison, but comparison theorists tend not to emphasize this dissimilarity (Tourangeau and Sternberg 1982, p. 205, their italics).

What is crucial in the comparison approach, then, is finding the correct ground in a metaphor. According to Tourangeau and Sternberg, Aristotle proposed the first

1 The met* method takes its name from a remark made by Yorick Wilks. He used met* to refer collectively to metonymy and metaphor: "*" is a match-anything symbol in the Unix operating system; hence, the token "met*" matches the two tokens "metonymy" and "metaphor."

comparison theory and suggested several principles for finding the ground of a metaphor. Tourangeau and Sternberg reduce these principles to two basic ones: finding a category to which the tenor and vehicle belong and constructing an analogy involving them.

Gentner's (1983) Structure-Mapping Theory, which has been implemented in the Structure-Mapping Engine (Falkenhainer, Forbus and Gentner 1989), closely resembles a comparison view of metaphor. The theory addresses literal similarity, analogy, abstraction, and anomaly, which Gentner refers to as four "kinds of comparison." An algorithm compares the semantic information from two concepts represented as sets of properties. Properties are either "attributes," one-place predicates like LARGE(x), or "relations," two-place predicates such as COLLIDE(x,y). The four kinds of comparison are distinguished by the relative proportions of attributes and relations that are matched, and the forms of mappings established between them. Mappings between relations are sought before those between attributes. Pairs of relations are compared using the "systematicity principle" that regular structural correspondences should exist between terms occupying the same positions in those relations. Mappings are purely structural and independent of the content of the relations (i.e., the predicates).

Tourangeau and Sternberg (1982) list some problems with the comparison view, including the following:

> (a) that everything has some feature or category that it shares with everything else, but we cannot combine just any two things in metaphor; (b) that the most obvious shared features are often irrelevant to a reading of the metaphor; (c) that even when the feature is relevant, it is often shared only metaphorically; ... and (e) that metaphors are novel and surprising is hard to reconcile with the idea that they rely completely on extant similarities (ibid., pp. 226–227).

Johnson (1980) also notes problem (a) with comparison theories, pointing out that as a result they cannot account for the semantic tension between the two terms of a metaphor:

> the comparison theory ... tries to circumvent the experienced semantic strain by interpreting metaphor as nothing but a way of comparing two things to see in what respects they are alike. And since any two things are similar in some respects, this kind of theory can never explain what is interesting and important about metaphor (ibid., p. 52).

2.1.2 The Interaction View. The interaction view focuses more upon the surprise and novelty that metaphors create. According to Tourangeau and Sternberg (1982, p. 212), proponents of the interaction view include Black (1962), Hesse (1966), Miles (1967), Richards (1936), and Wheelwright (1962).

> Interaction theorists argue that the vehicle of a metaphor is a template for seeing the tenor in a new way. This reorganization of the tenor is necessary, because the characteristics or features of the vehicle cannot be applied directly to the tenor; the features they 'share' are often only shared metaphorically. As Black (1962) observes, the ground of a metaphor may itself be nonliteral. 'Men are wolves,' in Black's example, in part because both are predators; but they are predators in sharply different senses that may only strike us as similar when we interpret the metaphor. In Black's reading of this metaphor, we see competition in social relations as corresponding to predacity in beasts (Tourangeau and Sternberg 1982, pp. 212–213).

A problem with the interaction view is that theorists have not provided much detail about the processes involved, though Black (1962) does make some suggestions.

> According to Black, tenor and vehicle... each have a 'system of commonplaces' associated with them. These commonplaces are stereotypes, not necessarily definitional, not even necessarily true, just widely agreed upon. In interpreting 'man is a wolf,' we 'evoke the wolf-system of related commonplaces' and are led by them 'to construct a corresponding system of implications about the principal subject (Man)' (Black, 1962, p. 41). In Black's view, then, interpretation involves not so much comparing tenor and vehicle for existing similarities, as construing them in a new way so as to create similarity between them (Tourangeau and Sternberg 1982, p. 213).

One might distinguish, then, two main differences between the interaction and comparison views. First, similarities are "created" in the interaction view (accounting for the novelty and surprise in a metaphor) whereas only pre-existing similarities are found in the comparison view. Second, a whole system of similarities are evoked between tenor and vehicle in the interactions view, whereas the comparisons view is based upon finding a single similarity.

One version of the interaction view is the domains-interaction view, set forth by Tourangeau and Sternberg (1982), who take the view that

> features 'shared' by tenor and vehicle are often at best only analogous features, each limited in its application to one domain or another. Of course, some features or dimensions are quite general, applying across the board to a number of domains (p. 218).

Among comparison and interaction theorists, much attention had been paid to selecting the comparisons or interactions in a metaphor. The importance of analogy or correspondence in metaphor has been stressed by Gentner (1983), Ortony (1979), Tourangeau and Sternberg (1982), and Wilks (1978), among others. Various mechanisms have been advanced for highlighting certain comparisons or interactions, including relevance (e.g., Hobbs 1983b; Tversky 1977) and salience (Ortony et al. 1985).

Among computational approaches, Indurkhya's (1988) Constrained Semantic Transference theory of metaphor can be viewed as a formalization of Black's interaction theory (ibid., p. 129). Source and target domains are viewed as "systems of relationships." In metaphorical interpretation, an "implicative complex" of the source domain is imposed on the target domain, thereby shaping the features of the target domain, which in turn produces changes in the features of the source domain, hence the "interaction." It is assumed that a structural analogy underlies every metaphor (ibid., p. 129).

A metaphor is identified with the formal notion of a T-MAP which is a pair $\langle F,S \rangle$ where F is a function that maps vocabulary of the source domain onto vocabulary of the target domain and S is a set of sentences from the source domain which are expected to transfer to the target domain. A metaphor is "coherent" if the transferred sentences S are logically consistent with the axioms of the target domain, and "strongly coherent" if they already lie in the deductive closure of those axioms (cf. Stallard 1987, p. 181). S is thus the "implicative complex" of the source domain imposed on the target domain. Every metaphorical interpretation of a given set of sentences is associated with a T-MAP. There may be several possible T-MAPs for a set of sentences.

I would argue that Hobbs (1983a; 1983b) has also taken an interaction view of metaphor. Hobbs' goal has been to develop a unified process of discourse interpretation based on the drawing of appropriate inferences from a large knowledge base, which Hobbs sometimes calls "selective inferencing" (e.g., Hobbs 1980). Selective inferencing is concerned with drawing or refraining from drawing certain inferences in a controlled fashion (cf. Hobbs 1983a). He argues that many problems have the same or almost the same inferencing solutions. These solutions are found via four separate semantic operations that all draw inferences from text (e.g., Hobbs 1977).

2.1.3 The Selection Restrictions Violations View. The selection restriction violation view has also been called "the semantic deviance view" (Johnson 1980, p. 50) and "the anomaly view" (Tourangeau and Sternberg 1982, p. 211). Johnson (1980) describes this view as a common one among linguists; Tourangeau and Sternberg (1982) list the following people as holders of this view: Beardsley (1962), Bickerton (1969), Campbell (1975), Guenther (1975), Percy (1954), Van Dijk (1975), and Wheelwright (1962). To this list one might add Levin (1977). Johnson (1980, p. 50) describes this view as where:

> metaphor constitutes a violation of selection restriction rules within a given context, where the fact of this violation is supposed to explain the semantic tension one experiences in comprehending any live metaphor.

The theory of metaphor in Preference Semantics (Wilks 1975; 1978) consists of a selection restrictions view and a comparison view. In the theory, information about word senses is contained in knowledge structures called "semantic formulas." An algorithm matches pairs of semantic formulas, seeking satisfied or violated preferences between them. A satisfied preference indicates a literal semantic relation; a violated preference indicates either a metaphorical or anomalous one. This part of the theory is implemented in a machine translation system (Wilks 1973).

To distinguish metaphor from anomaly, a different knowledge structure and a second algorithm are used. The algorithm, called projection, operates on a knowledge structure, called a pseudo-text, that contains lists of templates (a further kind of knowledge structure) linked by case ties. A brief example of projection is given for (1).

Example 3
"My car drinks gasoline."

Projection operates only on preference violations. The best representation of (1) contains a preference violation, so projection is used. The algorithm compares the template representation for the sentence

[my+car drink gasoline]

against templates from the pseudo-text of 'car' seeking "the closest match," and selects [ICengine (USE)#liquid]. (USE) is projected onto drink in the sentence representation which becomes

[my+car use gasoline]

Example 3
"The rock is becoming brittle with age" (Reddy 1969, p. 242).

Example 4
"Idi Amin is an animal" (Johnson 1980, p. 51).

Example 5
"People are not cattle" (Hobbs 1983b, p. 134).

Example 6
"No man is an Island" (John Donne, *Meditations XVII*).

The main problem with the selection restrictions view is that perfectly well-formed sentences exist that have a metaphorical interpretation and yet contain no selection restriction violations (Johnson 1980; Ortony 1980; Reddy 1969); for example, in (3), there is a literal interpretation when uttered about a stone and a metaphorical one when said about a decrepit professor emeritus. Sentences (4), (5) and (6) also have twin interpretations.

The existence of such sentences suggests that a condition that occasionally holds (i.e., a selection restriction violation) has been elevated into a necessary condition of metaphor (Johnson 1980). Moreover, viewing metaphor only in terms of selection restriction violations ignores the influence of context:

> We seem to interpret an utterance metaphorically when to do so makes sense of more aspects of the total context than if the sentence is read literally. Consider the simple case of the sentence *All men are animals* as uttered by Professor X to an introductory biology class and as uttered later by one of his female students to her roommate upon returning from a date. In the latter instance the roommate understands the utterance as metaphorical (ibid., p. 51).

In a similar way, Ortony (1980) suggests that metaphor should be thought of as

> *contextually* anomalous. This means that a literal interpretation of the expression, be it a word, phrase, sentence, or an even larger unit of text, fails to fit the context (p. 73, his italics),

so whether or not a sentence is a metaphor depends upon the context in which it is used:

> *if* something is a metaphor then it will be contextually anomalous if interpreted literally.... Insofar as the violation of selection restrictions can be interpreted in terms of semantic incompatibilities at the lexical level, such violations may sometimes be the basis of the contextual anomaly (ibid., p. 74).

2.1.4 The Conventional Metaphor View. Lakoff and Johnson (1980) have popularized the idea of conventional metaphors, also known as conceptual metaphors. They distinguish three main kinds: orientational, ontological, and structural. Orientational metaphors are mainly to do with kinds of spatial orientation like up-down, in-out, and deep-shallow. Example metaphors include MORE IS UP and HAPPY IS UP. They arise from human experience of spatial orientation and thus develop from the sort of bodies we have and the way they function in our physical environment.

Ontological metaphors arise from our basic human experiences with substances and physical objects (especially our own bodies). Some examples are TIME IS A SUBSTANCE, THE MIND IS AN ENTITY, and THE VISUAL FIELD IS A CONTAINER.

Structural metaphors are elaborated orientational and ontological metaphors (cf. Lakoff and Johnson 1980) in which concepts that correspond to natural kinds of experience, e.g., PHYSICAL ORIENTATIONS, SUBSTANCES, WAR, JOURNEYS, and BUILDINGS, are used to define other concepts, also natural kinds of experience, e.g., LOVE, TIME, IDEAS, UNDERSTANDING, and ARGUMENTS. Some examples of structural metaphors are ARGUMENT IS WAR and TIME IS MONEY.

> The ARGUMENT IS WAR metaphor forms a systematic way of talking about the battling aspects of arguing.... Because the metaphorical concept is systematic, the language we use to talk about the concept is systematic (ibid., p. 5).

What Lakoff and Johnson fail to discuss is how metaphors in general, let alone individual metaphorical concepts, are recognized. Martin's (1988; 1990) work has addressed this issue. He has pursued a conventional metaphor view using KODIAK (Wilensky 1984), a variant of Brachman's KLONE knowledge representation language. Within KODIAK, metaphorical relationships are represented using a primitive link type called a "VIEW." A VIEW "is used to assert that...one concept may in certain circumstances be considered as another " (Martin 1990, p. 59). In Martin's work, "metaphor-maps," a kind of VIEW (ibid., p. 64), are used to represent conventional metaphors and the conceptual information they contain.

2.2 Metonymy

Metonymy involves "using one entity to refer to another that is related to it" (Lakoff and Johnson 1980, p. 35).

Example 2
"The ham sandwich is waiting for his check."

For example, in (2) the metonymy is that the concept for ham sandwich is related to an aspect of another concept, for "the person who ordered the ham sandwich."

Several attempts have been made to organize instances of metonymy into categories (e.g., Lakoff and Johnson 1980; Stern 1931; Yamanashi 1987) or "metonymic concepts," as Lakoff and Johnson call them. A common metonymic concept is PART FOR WHOLE, otherwise known as synechdoche.

Example 7
"Dave drank the *glasses*" (= the liquid in the glasses).

Example 8
"The *kettle* is boiling" (= the liquid in the kettle) (Waldron 1967, p. 186; Yamanashi 1987, p. 78).

CONTAINER FOR CONTENTS, another metonymic concept, occurs in (7) between 'drink' and the sense of 'glasses' meaning "containers," and also in (8). In (7), 'drink' has an object preference for a potable liquid, but there is a preference violation because glasses are not potable liquids. It is not glasses that are drunk, but the potable *liquids* in them. There is a relationship here between a CONTAINER (a glass) and its typical CONTENTS (a liquid): this relationship is the metonymic concept CONTAINER FOR

CONTENTS. Below are examples of two further metonymic concepts (from Lakoff and Johnson 1980, p. 38, italics in original).

PRODUCER FOR PRODUCT
"I'll have a *Löwenbräu*."
"He bought a *Ford*."
"He's got a *Picasso* in his den."
"I hate to read *Heidegger*."

OBJECT USED FOR USER
"The *sax* has the flu today."
"The *BLT* is a lousy tipper."*[2]
"The *buses* are on strike."

Example 9
"You'll find better ideas than that in the library" (Reddy 1979, p. 309).

Reddy (1979) has observed that metonymies can occur in chains. He suggests that (9) contains a chain of PART FOR WHOLE metonymies between 'ideas' and 'library': the ideas are expressed in words, words are printed on pages, pages are in books, and books are found in a library.

Example 10
"I found an old car on the road. The steering wheel was broken" (Yamanashi 1987, p. 79).

Example 11
"We had a party in a mysterious room. The walls were painted in psychedelic color" (ibid.).

Example 12
A: "I bought an interesting book." B: "Who is the author?" (ibid.).

Example 13
"He happened to die of some disease, though I don't know what the cause was" (ibid.).

Yamanashi (1987) points out that basic metonymic relationships like part-whole and cause-result often also link sentences. According to him, the links in (10) and (11) are PART-WHOLE relations, the one in (12) is PRODUCT-PRODUCER, and the one in (13) is a CAUSE-RESULT relation.

There has been some computational work on metonymy (Weischedel and Sondheimer 1983; Grosz et al. 1987; Hobbs and Martin 1987; Stallard 1987; Wilensky 1987). The TEAM project (Grosz et al. 1987) handles metonymy, though metonymy is not mentioned by name but referred to instead as "coercion," which "occurs whenever some property of an object is used to refer indirectly to the object" (ibid., p. 213). Coercion is handled by "coercion-relations;" for example, a coercion relation could be used to understand that 'Fords' means "cars whose CAR-MANUFACTURER is Ford" (in Lakoff and Johnson's terms, this is an example of a PRODUCER FOR PRODUCT metonymic concept).

2 A BLT is a bacon, lettuce, and tomato sandwich.

Grosz et al. (1987) note a similarity between coercion (i.e., metonymy) and modification in noun-noun compounds, and use "modification relations" to decide whether, e.g., "U.S. ships" means "ships of U.S. registry" or "ships whose destination is the U.S."

Hobbs and Martin (1987) and Stallard (1987) also discuss the relationship between metonymy and nominal compounds. Hobbs and Martin treat the two phenomena as twin problems of reference resolution in their TACITUS system. They argue that resolving reference requires finding a knowledge base entity for an entity mentioned in discourse (i.e., what that entity refers to), and suggest that the resolution of metonymy and nominal compounds both require discovering an implicit relation between two entities referred to in discourse. The example of metonymy they show is "after the alarm," which really means after the *sounding* of the alarm.

Hobbs and Martin seem to assume a selection restrictions approach to metonymy because metonymy is sought after a selection restrictions violation (ibid., p. 521). In their approach, solving metonymy involves finding: [1] the referents for 'after' and 'alarm' in the domain model, which are *after*(e_0, a) and *alarm*(a); [2] an implicit entity z to which 'after' really refers, which is *after*(e_0, z); and [3] the implicit relation between the implicit entity z and the referent of 'alarm,' q(z, a).

Like Hobbs and Martin (1987), Stallard (1987) translates language into logical form. Stallard argues that with nominal compounds and metonymies "the problem is determining the binary relation which has been 'elided' from the utterance" (ibid., p. 180) and suggests shifting the argument place of a predicate "by interposing an arbitrary, sortally compatible relation between an argument place of the predicate and the actual argument" (ibid., p. 182). Stallard notes that "in any usage of the metonomy (sic) operation there is a choice about which of two clashing elements to extend" (ibid.). Stallard's work has not yet been implemented (ibid., p. 184).

Stallard (1987) also briefly discusses anaphora resolution. Brown (1990) is beginning research on metonymy and reference resolution, particularly pronouns. This should prove a promising line of investigation because metonymy and anaphora share the function of allowing one entity to refer to another entity.

Example 2
"The ham sandwich is waiting for his check" (= the male person who ordered the ham sandwich).

Example 14
"He is waiting for his check" (= the male person).

This similarity of function can be seen in comparing (2), which is metonymic, with (14), which is anaphoric.

2.3 Relationship between Metonymy and Metaphor

Both metonymy and metaphor have been identified as central to the development of new word senses, and hence to language change (see, e.g., Stern 1931; Waldron 1967). Some of the best examples of the differences between the two phenomena come from data used in studies of metonymic and metaphorical effects on language change. Nevertheless, there are widely differing views on which phenomenon is the more important. Some argue that metaphor is a kind of metonymy, and others propose that metonymy is a kind of metaphor, while still others suggest that they are quite different (see Fass 1988c).

Among the third group, two differences between metonymy and metaphor are commonly mentioned. One difference is that metonymy is founded on contiguity whereas metaphor is based on similarity (cf. Jakobsen and Halle 1956; Ullmann 1962). Contiguity and similarity are two kinds of association. Contiguity refers to a state of being connected or touching whereas similarity refers to a state of being alike in essentials or having characteristics in common (Mish 1986).

A second difference, advanced by Lakoff and Johnson (1980) for example, is that metaphor is "principally a way of conceiving of one thing in terms of another, and its primary function is understanding" (ibid., pp. 36–37) whereas metonymy "has primarily a referential function, that is, it allows us to use one entity to *stand for* another" (ibid., their italics), though it has a role in understanding because it focuses on certain aspects of what is being referred to.

There is little computational work about the relationship between metonymy and metaphor. Stallard (1987) distinguishes separate roles for metonymy and metaphor in word sense extension. According to him, metonymy shifts the argument place of a predicate, whereas metaphor shifts the whole predicate. Hobbs (1983a; 1983b) writes about metaphor, and he and Martin (1987) develop a theory of "local pragmatics" that includes metonymy, but Hobbs does not seem to have written about the relationship between metaphor and metonymy.

In knowledge representation, metonymic and metaphorical relations are both represented in the knowledge representation language CycL (Lenat and Guha 1990).

2.4 Literalness and Nonliteralness

Much of the preceding material assumes what Gibbs (1984) calls the "literal meanings hypothesis," which is that

> sentences have well defined literal meanings and that computation of the literal meaning is a necessary step on the path to understanding speakers' utterances (ibid., p. 275).

There are a number of points here, which Gibbs expands upon in his paper. One point concerns the traditional notion of literal meaning, that all sentences have literal meanings that are entirely determined by the meanings of their component words, and that the literal meaning of a sentence is its meaning independent of context.

A second point concerns the traditional view of metaphor interpretation, though Gibbs' criticism applies to metonymy interpretation also. Using Searle's (1979) views on metaphor as an example, he characterizes the typical model for detecting nonliteral meaning as a three-stage process: [1] compute the literal meaning of a sentence, [2] decide if the literal meaning is defective, and if so, [3] seek an alternative meaning, i.e., a metaphorical one (though, presumably, a metonymic interpretation might also be sought at this stage). Gibbs (1984, p. 275) concludes that the distinction between literal and metaphoric meanings has "little psychological validity."

Among AI researchers, Martin (1990) shares many of Gibbs's views in criticizing the "literal meaning first approach" (ibid., p. 24). Martin suggests a two-stage process for interpreting sentences containing metaphors: [1] parse the sentence to produce a syntactic parse tree plus primal (semantic) representation, and [2] apply inference processes of "concretion" and "metaphoric viewing" to produce the most detailed semantic representation possible.

> The primal representation represents a level of semantic interpretation that is explicitly in need of further processing. Although it is obviously related to what

> has traditionally been called a literal meaning, it should not be thought of as a *meaning* at all. The primal representation should be simply considered as an intermediate stage in the interpretation process where only syntactic and lexical information has been utilized (ibid., p. 90, his italics).

However, Martin believes that at least some sentence meaning is independent of context because the primal representation contains part of the primal content of an utterance and

> [t]he Primal Content represents the meaning of an utterance that is derivable from knowledge of the conventions of a language, independent of context (ibid.).

2.5 Review Summary

The metaphor literature contains many differing views, including the comparison, interaction, selection restrictions, and conventional metaphors views. AI research on metaphor includes all of these views. Of the AI research, only Martin's work has been implemented to my knowledge. Among the points raised are that metaphorical sentences exist that do not contain selection restriction violations and that metaphor requires interpretation in *context*. The much smaller metonymy literature stresses the selection restrictions view too. The TEAM and TACITUS systems both seem to process metonymics.

The two main differences commonly noted between metonymy and metaphor are in their function (referential for metonymy and understanding with metaphor) and the kind of relationship established (contiguity in metonymy versus similarity in metaphor). No one to my knowledge has a working system that discriminates examples of metaphor and metonymy.

3. met* Method

In this section, the basic met* algorithm is outlined. The met* method is based on the selection restriction, also known as the preference. Metonymy, metaphor, literalness, and anomaly are recognized by evaluating preferences, which produces four kinds of basic "preference-based" relationship or *semantic relation*: literal, metonymic, metaphorical, and anomalous. Within the method, the main difference between metonymy and metaphor is that a metonymy is viewed as consisting of one or more semantic relationships like CONTAINER FOR CONTENTS and PART FOR WHOLE, whereas a metaphor is viewed as containing a relevant analogy.

I agree with Ortony's remark that metaphor be viewed as contextual anomaly, but would suggest two modifications. First, not just metaphor but all of the preference-based relations should be understood in terms of the presence or absence of contextual constraint violation. Second, I prefer the term *contextual constraint violation* because [1] one of the phenomena detected by contextual violation *is* anomaly and [2] the selection restriction/preference (on which the met* method is based) is a kind of lexical contextual constraint. The section starts with an explanation of some of the linguistic background behind the met* method.

3.1 Linguistic Background

I have argued elsewhere (Fass 1989a) that understanding natural language (or semantic interpretation) be viewed as the integration of *constraints* from language and from context. Some language constraints are syntactic, while others are semantic. Some language constraints are *lexical* constraints; that is, constraints possessed by lexical items (words and fixed phrases). Lexical syntactic constraints include those on word order, number, and tense. This section describes three *lexical semantic constraints: preferences, assertions,* and a lexical notion of *relevance.*

Preferences (Wilks 1973), selection restrictions (Katz 1964), and expectations (Schank 1975) are the same (see Fass 1989c; Fass and Wilks 1983; Wilks and Fass in press): all are restrictions possessed by senses of lexical items of certain parts of speech about the semantic classes of lexical items with which they co-occur. Thus an adjective sense has a preference for the semantic class of nouns with which it co-occurs and a verb sense has preferences for the semantic classes of nouns that fill its case roles. For example, the main sense of the verb 'drink' prefers an animal to fill its agent case role, i.e., it is animals that drink.

The assertion of semantic information was noted by Lees (1960) in the formation of noun phrases and later developed by Katz (1964) as the process of "attribution." Assertions contain information that is possessed by senses of lexical items of certain parts of speech and that is imposed onto senses of lexical items of other parts of speech, e.g., the adjective 'female' contains information that any noun to which it applies is of the female sex.

Lexical syntactic and semantic constraints are enforced at certain places in sentences which I call *dependencies.* Within a dependency, the lexical item whose constraints are enforced is called the *source* and the other lexical item is called the *target* (after Martin 1985). Syntactic dependencies consist of pairs of lexical items of certain parts of speech in which the source, an item from one part of speech, applies one or more syntactic constraints to the target, another lexical item. Examples of source-target pairs include a determiner and a noun, an adjective and a noun, a noun and a verb, and an adverb and a verb.

Example 15
"The ship ploughed the waves."

Semantic dependencies occur in the same places as syntactic dependencies. The (metaphorical) sentence (15) contains four semantic dependencies: between the determiner 'the' and the noun 'ship,' between 'ship' and the verb stem 'plough,' between 'the' and the noun 'waves,' and between 'waves' and 'plough.' In each semantic dependency, one lexical item acts as the source and applies constraints upon the other lexical item, which acts as the target. In (15), 'the' and 'plough' both apply constraints upon 'ship,' and 'the' and 'plough' apply constraints on 'waves.' Semantic dependencies exist between not just pairs of lexical items but also pairs of *senses* of lexical items. For example, the metaphorical reading of (15) is because 'waves' is understood as being the sense meaning "movement of water," not for example the sense meaning "movement of the hand."

Semantic relations result from evaluating lexical semantic constraints in sentences. Every semantic relation has a source (a lexical item whose semantic constraints are applied) and a target (a lexical item which receives those constraints). Other terms used to refer to the source and target in a semantic relation include: vehicle and tenor (Richards 1936), subsidiary subject and principal subject (Black 1962), figurative term and literal term (Perrine 1971), referent and subject (Tversky 1977), secondary subject and primary subject (Black 1979), source and destination (Winston 1980), old domain and new domain (Hobbs 1983a), and base and target (Gentner 1983).

In CS, seven kinds of semantic relation are distinguished: literal, metonymic, metaphorical, anomalous, redundant, inconsistent, and novel relations (this list may not be exhaustive — there could be others). Combinations of these seven semantic relations are the basis of (at minimum) literalness, metonymy, metaphor, anomaly, redundancy, contradiction, contrariness, and novelty. Semantic relations belong to two classes, the *preference-based* and *assertion-based* classes of relations, depending on the kind of lexical semantic constraint enforced. The preference-based class of semantic relations, which are the focus of this paper, contains literal, metonymic, metaphorical, and anomalous semantic relations. The assertion-based class of relations are described in greater length in Fass (1989a).

Figure 1 shows the met* method laid out as a flow chart and illustrates how the preference-based class of semantic relations is discriminated. A satisfied preference (diamond 1) distinguishes literal relations from the remaining three relations, which are all nonliteral.

Example 16
"The man drank beer."

There is a literal relation between 'man' and 'drink' in (16) because 'drink' prefers an animal as its agent and a man is a type of animal so the preference is satisfied.

Example 7
"Dave drank the *glasses*" (= potable liquid in the glasses → CONTAINER FOR CONTENTS).

Example 17
"Denise drank the *bottle*" (= potable liquid from the bottle → CONTAINER FOR CONTENTS).

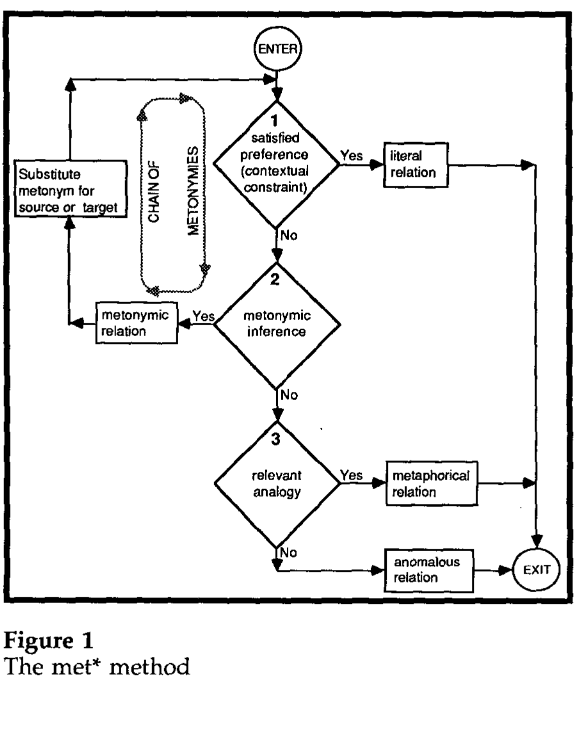

Figure 1
The met* method

Example 18
"Anne reads *Steinbeck*" (= writings of Steinbeck → ARTIST FOR ART FORM).

Example 19
"Ted played *Bach*" (= music of Bach → ARTIST FOR ART FORM).

Metonymy is viewed as a kind of domain-dependent inference. The process of finding metonymies is called *metonymic inferencing*. The metonymic concepts presently used are adapted from the metonymic concepts of Lakoff and Johnson (1980). Two of the metonymic concepts used are CONTAINER FOR CONTENTS and ARTIST FOR ART FORM. In (19), for example, Ted does not literally play the composer Bach — he plays music composed by him.

As Figure 1 shows, a metonymy is recognized in the met* method if a metonymic inference (diamond 2) is found. Conversely, if no successful inference is found then no metonymy is discovered and a metaphorical or anomalous semantic relation is then sought. A successful inference establishes a relationship between the original source or the target ("one entity") and a term ("another that is related to it") that refers to one of them.

Like Stallard (1987), who noted that "in any usage of the metonomy (sic) operation there is a choice about which of two clashing elements to extend" (ibid., p. 182), the met* method allows for metonymies that develop in different "directions." A successful inference is sometimes directed "forward" from the preference or "backward" from the target, depending on the metonymic concept (more on this shortly). It is this direction of inferencing that determines whether the source or target is substituted in a successful metonymy. The substitute source or target is used to discover another semantic relation that can be literal, metonymic again, metaphorical, or anomalous.

In Figure 1, the presence of a relevant analogy (diamond 3) discriminates metaphorical relations from anomalous ones. No one else (to my knowledge) has emphasized the role of relevance in the discovery of an analogy central to a metaphor though, as noted in Section 2.2, the importance of relevance in recognizing metaphors and the centrality of some analogy have both been discussed.

Example 20
"The car drank gasoline" (adapted from Wilks 1978).

The form of relevance used is a lexical notion — i.e., the third kind of lexical semantic constraint — that what is relevant in a sentence is given by the sense of the main sentence verb being currently analyzed. Thus, it is claimed that the semantic relation between 'car' and 'drink' in (20) is metaphorical because there is a preference violation and an underlying relevant analogy between 'car' and 'animal,' the preferred agent of 'drink.' A car is not a type of animal, hence the preference violation. However, what is relevant in (20) is drinking, and there is a relevant analogy that animals and cars both use up a liquid of some kind: animals drink potable liquids while cars use gasoline. Hence the metaphorical relation between 'car' and 'drink.'

Metaphor recognition in the met* method is related to all four views of metaphor described in Section 2. Recognition is viewed as a two-part process consisting of [1] a contextual constraint violation and [2] a set of "correspondences" including a key correspondence, a relevant analogy. The contextual constraint violation may be a preference violation, as in the selection restrictions view of metaphor. The set of "correspondences" is rather like the system of commonplaces between tenor and vehicle in the interaction view. The relevant analogy is related to the comparison and interaction views, which emphasize a special comparison or an analogy as central to metaphor. Moreover, the relevant analogies seem to form groupings not unlike the conceptual metaphors found in the conventional view.

Example 21
"The idea drank the heart."

Anomalous relations have neither the semantic relationships of a metonymic relation nor the relevant analogy of a metaphorical relation. Hence the semantic relation between 'idea' and 'drink' is anomalous in (21) because 'idea' is not a preferred agent of 'drink' and no metonymic link or relevant analogy can be found between animals (the preferred agent) and ideas; that is, 'idea' in (21) does not use up a liquid like 'car' does in (20). This is not to say that an anomalous relation is uninterpretable or that no analogy can possibly be found in one. In special circumstances (for example, in a poem), search for analogies might be expanded to permit weaker analogies, thereby allowing "ideas drinking" to be interpreted metaphorically.

The topology of the flow chart in Figure 1 results from needing to satisfy a number of observations about the preference-based phenomena, particularly metonymy:

1. literalness is distinct from the others, which are all nonliteral;
2. metonymies can occur in chains (Reddy 1979);
3. metonymy always seems to occur with one of the other three; and
4. metaphor and anomaly are the hardest to tell apart (and thus require the most extended processing to distinguish).

Hence a preference-based semantic relation can be either a single relation or a *multi-relation*. A single relation consists of one literal, metaphorical, or anomalous relation. A multi-relation contains one literal, metaphorical, or anomalous relation plus either a single metonymy or a chain of metonymies. All these combinations, but only these, are derivable from Figure 1.

Note that in the met* method as presented in Figure 1, semantic relations are tried in a certain order: literal, metonymic, metaphorical, and finally anomalous. This ordering implies that a literal interpretation is sought *before* a nonliteral one (cf. Harris 1976). The ordering results from thinking about discriminating the semantic relations in serial processing terms rather than parallel processing terms, particularly the serial order in which selection restrictions are evaluated and metonymic inference rules are tried: satisfied selection restrictions (indicating literalness) then metonymic inference (metonymy) then violated selection restrictions (metaphor or anomaly).

Gibbs (1984) criticizes the idea that literal and nonliteral meaning can be discriminated in ordered processing stages. My response is that if the met* method is viewed in parallel processing terms then literal, metonymic, metaphorical, and anomalous interpretations are all sought at the same time and there is no ordering such that the literal meaning of a sentence is computed first and then an alternative meaning sought if the literal meaning is defective. Gibbs' other main criticism, concerning the traditional analysis of sentence meaning as composed from word meanings and independent of context, will be discussed in Section 7.

4. Collative Semantics

CS is a semantics for natural language processing that extends many of the main ideas behind Preference Semantics (Wilks 1973; 1975a; 1975b; 1978; see also Wilks and Fass in press). CS has four components: *sense-frames, collation, semantic vectors,* and *screening*. The met* method is part of the process of collation. Fuller and more general descriptions of the four components appear in Fass (1988a; 1989b).

Sense-frames are dictionary entries for individual word senses. Sense-frames are composed of other word senses that have their own sense-frames, much like Quillian's (1967) planes. Each sense-frame consists of two parts, an *arcs* section and a *node* section, that correspond to the genus and differentia commonly found in dictionary definitions (Amsler 1980).

The arcs part of a sense-frame contains a labeled arc to its genus term (a word sense with its own sense-frame). Together, the arcs of all the sense-frames comprise a densely structured semantic network of word senses called the *sense-network*. The node part of a sense-frame contains the differentia of the word sense defined by that sense-frame, i.e., information distinguishing that word sense from other word senses sharing the same genus. The two lexical semantic constraints mentioned earlier, preferences and assertions, play a prominent part in sense-frame nodes.

Sense-frame nodes for nouns (node-type 0) resemble Wilks' (1978) pseudo-texts. The nodes contain lists of two-element and three-element lists called *cells*. Cells contain word senses and have a syntax modeled on English. Each cell expresses a piece of functional or structural information and can be thought of as a complex semantic feature or property of a noun. Figure 2 shows sense-frames for two senses of the noun 'crook.' Crook1 is the sense meaning "thief" and crook2 is the shepherd's tool.

All the terms in sense-frames are word senses with their own sense-frames or words used in a particular sense that could be replaced by word senses. It1 refers to the word sense being defined by the sense-frame so, for example, crook1 can be substituted for it1 in [it1, steal1, valuables1]. Common dictionary practice is followed in that word senses are listed separately for each part of speech and numbered by frequency of occurrence. Hence in crook2, the cell [shepherd1, use1, it1] contains the noun sense shepherd1 while the cell [it1, shepherd1, sheep1] contains the verb sense shepherd1 (in a three-element cell, the second position is always a verb, and the first and third positions are always nouns).

Sense-frame nodes for adjectives, adverbs and other modifiers (node-type 1) contain preferences and assertions but space does not permit a description of them here. Sense-frame nodes for verbs and prepositions (node-type 2) are case frames containing case subparts filled by case roles such as 'agent,' 'object,' and 'instrument.' Case subparts contain preferences, and assertions if the verb describes a state change.

```
sf(crook1,                                   sf(crook2,
   [[arcs,                                      [[arcs,
     [[supertype,   criminal1]]],                 [[supertype,   stick1]]],
    [node0,                                      [node0,
     [[it1,   steal1,   valuables1]]]]).          [[shepherd1,   use1,   it1],
                                                   [it1,   shepherd1,   sheep1]]]]).
```

Figure 2
Sense-frames for crook1 and crook2 (noun senses)

```
sf(eat1,
  [[arcs,
    [[supertype, [ingest1, expend1]]]],
   [node2,
    [[agent,
      [preference, animal1]],
     [object,
      [preference, food1]]]]]).
```

```
sf(drink1,
  [[arcs,
    [[supertype, [ingest1, expend1]]]],
   [node2,
    [[agent,
      [preference, animal1]],
     [object,
      [preference, drink1]]]]]).
```

Figure 3
Sense-frames for eat1 and drink1 (verb senses)

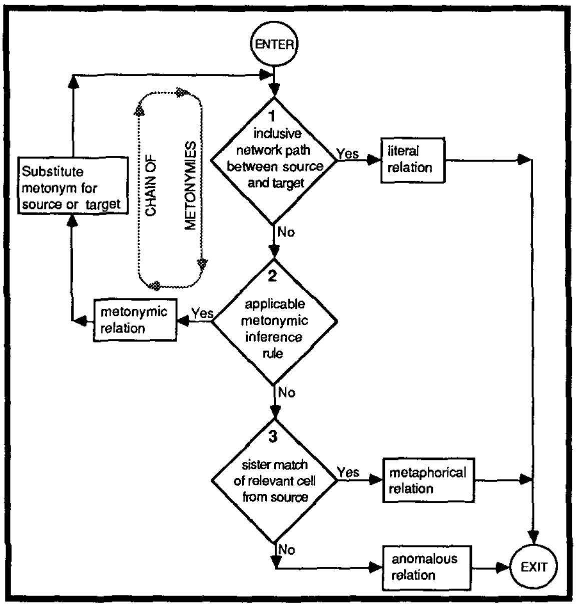

Figure 4
The met* method (CS version)

Figure 3 shows the sense-frames for the verb senses eat1 and drink1. In both, the agent preference is for an animal but the object preferences differ: the preference of eat1 is for food1, i.e., an edible solid, while the preference of drink1 is for drink1 (the noun sense), i.e., a potable liquid.

The second component of CS is the process of collation. It is collation that contains the met* method in CS. Collation matches the sense-frames of two word senses and finds a system of multiple mappings between those sense-frames, thereby discriminating the semantic relations between the word senses. Figure 4 shows the use of the met* method in CS. Figure 4 is similar to the one in Figure 1 except that the diamonds contain the processes used in CS to check for satisfied preferences (diamond 1), metonymic inferences (diamond 2), and relevant analogies (diamond 3).

The basic mappings in collation are paths found by a graph search algorithm that operates over the sense-network. Five types of network path are distinguished. Two types of path, called *ancestor* and *same*, denote kinds of "inclusion," e.g., that the class of vehicles includes the class of cars (this is an ancestor relationship). Satisfied preferences are indicated by network paths denoting inclusion, also known as "inclusive" paths (see diamond 1 in Figure 4). The other three types of network path, called *sister, descendant,* and *estranged,* denote "exclusion," e.g., that the class of cars does not include the class of vehicles (this is a descendant relationship). Violated preferences are network paths denoting exclusion, also known as "exclusive" paths.

These paths are used to build more complex mappings found by a frame-matching algorithm. The frame-matching algorithm matches the sets of cells from two sense-frames. The sets of cells, which need not be ordered, are inherited down the sense-network. A series of structural constraints isolate pairs of cells that are matched using the graph search algorithm. Network paths are then sought between terms occupying identical positions in those cells. Seven kinds of cell match are distinguished, based on the structural constraints and types of network path found. *Ancestor* and *same* are "inclusive" cell matches, e.g. [composition1, metal1] includes [composition1, steel1] because the class of metals includes the class of steels (another ancestor relationship). *Sister, descendant,* and *estranged* are types of "exclusive" cell matches, e.g. [composition1, steel1] and [composition1, aluminium1] are exclusive because the class of steels does not include the class of aluminiums since both belong to the class of metals (this is a sister relationship). The remaining cell matches, *distinctive source* and *distinctive target,* account for cells that fail the previous five kinds of cell match. For more detail on cell matches, see Fass (1988a).

A kind of lexical relevance is found dynamically from the sentence context. This notion of relevance is used in finding the relevant analogies that distinguish metaphorical from anomalous relations; it is also used when finding CO-AGENT FOR ACTIVITY metonymies. Relevance divides the set of cells from the source sense-frame into two subsets. One cell is selected as *relevant* given the context; the remaining cells are termed *nonrelevant.* Collation matches both the source's relevant and nonrelevant cells against the cells from the target sense-frame. A relevant analogy is indicated by a sister match of the source's relevant cell (see diamond 3 in Figure 4).

Five types of metonymic concepts are currently distinguished. Examples of two of the metonymic concepts, CONTAINER FOR CONTENTS and ARTIST FOR ART FORM, have already been given. The remaining three are PART FOR WHOLE, PROPERTY FOR WHOLE, and CO-AGENT FOR ACTIVITY.

Example 22
"Arthur Ashe is *black*" (= skin colored black → PART FOR WHOLE).

Example 23
"John McEnroe is *white*" (= skin colored white → PART FOR WHOLE).

In (22) and (23), the skins of Arthur Ashe and John McEnroe, parts of their bodies, are colored black (white).

Example 24
"John McEnroe is *yellow*" (= limited in bravery → PROPERTY FOR WHOLE).

Example 25
"Natalia Zvereva is *green*" (= limited in experience → PROPERTY FOR WHOLE).

In (24), for example, John McEnroe is limited with respect to his bravery, a property possessed by humans and other animals.

Example 26
"Ashe played *McEnroe*" (= tennis with McEnroe → CO-AGENT FOR ACTIVITY).

These concepts are encoded in *metonymic inference rules* in CS (see diamond 2 in Figure 4). The rules are ordered from most common (synecdoche) to least. The order used is PART FOR WHOLE, PROPERTY FOR WHOLE, CONTAINER FOR CONTENTS, CO-AGENT FOR ACTIVITY, and ARTIST FOR ART FORM.

The first two concepts, PART FOR WHOLE and PROPERTY FOR WHOLE, are source-driven; the others are target-driven. The difference in direction seems to be dependent on the epistemological structure of the knowledge being related by the different inferences. PART FOR WHOLE metonymies are source-driven, perhaps because the epistemological nature of parts and wholes is that a part generally belongs to fewer wholes than wholes have parts, hence it makes sense to drive inferencing from a part (source) toward the whole (target) than vice versa.

In CONTAINER FOR CONTENTS (target-driven), on the other hand, the epistemological nature of containers and contents is that the containers generally mentioned in CONTAINER FOR CONTENTS metonymies are artifacts designed for the function of containing — hence one can usually find quite specific information about the typical contents of a certain container, for example, some glasses as in (7) — whereas the contents do not generally have the function of being the contents of something. Hence it makes sense to drive inferencing from the container, and the function it performs, toward the contents than vice versa. The same reasoning applies to ARTIST FOR ART FORM (target-driven). An artist has the vocation of creating art: that is his/her purpose.

A further step in collation distinguishes metaphorical from anomalous semantic relations. Recall that a metaphorical relation contains a relevant analogy, as in (15) and (20), while an anomalous relation does not, as in (21). A relevant analogy is found by matching the relevant cell from the source sense-frame with one of the cells from the target sense-frame. If the match of cells is composed of a set of sister network paths between corresponding word senses in those cells, then this is interpreted as analogical and hence indicative of a metaphorical relation. Any other match of cells is interpreted as not analogical and thus an anomalous semantic relation is recognized (see Fass 1986; 1987).

The third component of CS is the semantic vector which is a form of representation, like the sense-frame; but sense-frames represent lexical knowledge, whereas semantic vectors represent coherence. Semantic vectors are therefore described as a kind of *coherence representation*. A semantic vector is a data structure that contains nested labels and ordered arrays structured by a simple dependency syntax. The labels form into sets. The outer sets of labels indicate the application of the three kinds of lexical semantic constraints. The outermost set of labels is 'preference' and 'assertion.' The middle set is 'relevant' and 'nonrelevant.' The innermost set is the kind of mapping used: 'network path' and 'cell matches.' The nesting of labels shows the order in which each source of knowledge was introduced. The ordered arrays represent the subkinds of each kind of mapping. Five-column arrays are for the five network paths; seven-column arrays are for the seven types of cell match. Each column contains a positive number that shows the number of occurrences of a particular network path or cell match.

The fourth component of CS is the process of screening. During analysis of a sentence constituent, a semantic vector is created for every pairwise combination of word senses. These word sense combinations are called *semantic readings* or simply "readings." Each reading has an associated semantic vector. Screening chooses between two semantic vectors and hence their attached semantic readings. Rank orderings among semantic relations are applied. In the event of a tie, a measure of conceptual similarity is used.

The ranking of semantic relations aims to achieve the most coherent possible interpretation of a reading. The class of preference-based semantic relations takes precedence over the class of assertion-based semantic relations for lexical disambiguation. The rank order among preference-based semantic relations is

$$\text{literal} \rightarrow \text{metaphorical} \rightarrow \text{anomalous.}$$

If the semantic vectors are still tied then the measure of conceptual similarity is employed. This measure was initially developed to test a claim by Tourangeau and Sternberg (1982) about the aptness of a metaphor. They contend that aptness is a function of the distance between the conceptual domains of the source and target involved: the claim is that the more distant the domains, the better the metaphor. This is discussed further in Section 5. The conceptual similarity measure is also used for lexical ambiguity resolution (see Fass 1988c).

5. The Meta5 Program

CS has been implemented in the meta5 natural language program. The meta5 program is written in Quintus Prolog and consists of a lexicon holding the sense-frames of just over 500 word senses, a small grammar, and semantic routines that embody collation and screening, the two processes of CS. The program is syntax-driven, a form of control carried over from the structure of earlier programs by Boguraev (1979) and Huang (1985), on which meta5 is based. Meta5 analyzes sentences, discriminates the seven kinds of semantic relation between pairs of word senses in those sentences (i.e., the program recognizes metonymies, metaphors, and so on), and resolves any lexical ambiguity in those sentences. Meta5 analyzes all the sentences given in Sections 3 and 4, plus a couple more metaphorical sentences discussed in Section 7.

Below are simplified versions of some of the metonymic inference rules used in meta5. The metonymic concepts used in CS contain three key elements: the conceptual relationship involved, the direction of inference, and a replacement of the source or target. The metonymic inference rules in meta5 contain all three key elements. The rules, though written in a prolog-like format, assume no knowledge of Prolog on the part of the reader and fit with the role of metonymy shown in Figures 1 and 4.

Each metonymic inference rule has a left-hand side and a right-hand side. The left-hand side is the topmost statement and is of the form *metonymic_inference_rule*(Source, Target). The right-hand side consists of the remaining statements. These statements represent the conceptual relationship and the direction of inference, except for the bottom most one, which controls the substitution of the discovered metonym for either the source or target: this statement is always a call to find a new sense-network path.

Rule 1
PROPERTY FOR WHOLE: source-driven.
metonymic_inference_rule (Source, Target):-
    find_cell (Source, [Whole, have1, it1]), [1]
    find_sense_network_path (Whole, Target). [2]

This rule represents PROPERTY FOR WHOLE, which is source-driven. Statement [1] represents the conceptual relationship and direction of inference. The conceptual relationship is that the source is a property possessed by the whole in a property-whole relation. The inference is driven from the source: *find_cell* searches through the source's list of cells for one referring to a "whole" of which the source is a "part." Statement [2] controls the substitution of the discovered metonym: the "whole" is the substitute metonym that replaces the source, and the next sense-network path is sought between the whole and the target.

Rule 2
CONTAINER FOR CONTENTS: target-driven.
metonymic_inference_rule (Source, Target):-
find_cell (Target, [it1, contain1, Contents]), [1]
find_sense_network_path (Source, Contents). [2]

This metonymic concept is target-driven. The target is the "container" in a container-contents relation ([1]). The "contents" is the substitute metonym that replaces the target. The next sense-network path is sought between the source and the contents ([2]).

Rule 3
ARTIST FOR ART FORM: target-driven.
metonymic_inference_rule (Source, Target):-
find_genus (Target, Occupation), [1]
find_cell (Occupation, [it1, Make, Art form]), [2]
confirm_type (create1, Make), [3]
confirm_type (art_form1, Art form), [4]
find_sense_network_path (Source, Art form). [5]

Again, the inference in ARTIST FOR ART FORM is from the target. The target is a person who is an "artist" in an artist–art form relation. The occupation of the person is found by searching up the sense-network ([1]). The list of cells associated with the occupation are searched for a cell describing the main activity involved in the occupation ([2]), e.g., a cook cooks food and an artist makes art forms. Checks are done to confirm that any activity found is indeed making an art form, i.e., that the "making" involved is a type of creating ([3]) and that the "art form" is a type of art form1 ([4]). The "art form" is the substitute metonym that replaces the target. A new sense-network path is computed between the source and the art form ([5]). I will now describe how meta5 recognizes some metonymies and metaphors.

Example 19
"Ted played *Bach*" (= the music of Bach).

In (19), between 'Bach' and the twelfth sense of 'play' in meta5's lexicon (meaning "to play music"), there is a chain of metonymies plus a literal relation. The chain consists of ARTIST FOR ART FORM and CONTAINER FOR CONTENTS metonymies. Both metonymic concepts are target-driven. In ARTIST FOR ART FORM the inference is from the ARTIST (the target) to the ART FORM (the source), so the substitute metonym replaces the target (the ARTIST) if the inference is successful.

The sense-frames of the verb sense play12 and the noun senses music1 and johann_sebastian_bach are shown in Figure 5. The semantic relation results from matching the object preference of play12, which is for music, against the surface object, which is 'Bach,' short for 'Johann Sebastian Bach.' The preference is the source and the surface object is the target.

We will follow what happens using the flow chart of Figure 4. (Enter diamond 1 of the chart.) The sense-network path between the source (music1) and the target

```
sf(play12,
   [[arcs,
     [[supertype, perform1]]],
    [node2,
      [[agent,
        [preference, human_being1]],
       [object,
        [preference,music1]]]]]).

sf(music1,
   [[arcs,
     [[supertype, [sound1, art_form1]]]],
    [node0,
     [[musician1, play12, it1]]]]).

sf(johann_sebastian_bach,
   [[arcs,
     [[supertype, composer1]]],
    [node0,
     [[animacy1, dead1],
      [sex1, male1],
      [born1, 1685],
      [died1, 1750]]]]).
```

Figure 5
Sense-frames for play12 (verb sense), music1 and johann_sebastian_bach (noun senses)

(johann_sebastian_bach) is sought. The path is not inclusive because johann_sebastian_bach is not a type of music1.

(Enter diamond 2 of the chart.) Metonymic inference rules are applied. The rules for PART FOR WHOLE, PROPERTY FOR WHOLE, CONTAINER FOR CONTENTS, CO-AGENT FOR ACTIVITY are tried in turn, but all fail. The rule for ARTIST FOR ART FORM, however, succeeds. The discovered metonymic inference is that johann_sebastian_bach (the ARTIST) composes musical pieces (the ART FORM). The metonymic inference is driven from the target (the ARTIST), which is johann_sebastian_bach. The successful metonymic inference, using the ARTIST FOR ART FORM inference rule above, is as follows: [1] johann_sebastian_bach (the ARTIST) is a composer1, [2] composers compose1 musical pieces (the ART FORM). Additional tests confirm [2], which are that [3] composing is a type of creating, and [4] a musical_piece1 is a type of art_form1.

(Enter the leftmost statement box — also step [5] of the ARTIST FOR ART FORM inference rule above.) The original target (johann_sebastian_bach) is replaced by the substitute metonym (musical_piece1).

(Enter diamond 1 for a second time.) The sense-network path between the source (music1) and the new target (musical_piece1) is sought. The path is not inclusive.

(Enter diamond 2 for a second time.) Metonymic inference rules are applied. The rules for PART FOR WHOLE and PROPERTY FOR WHOLE fail, but the rule for CONTAINER FOR CONTENTS succeeds. The successful inference, using the description of the CONTAINER–CONTENTS inference rule given previously, is that [1] a musical_piece1 (the CONTAINER) contains music1 (the CONTENTS).

(Enter the leftmost statement box for a second time.) The direction of inference in the CONTAINER FOR CONTENTS metonymic concept is from the target (the CONTAINER) towards the source (the CONTENTS), so [2] the target (the CONTAINER) is replaced by the substitute metonym when an inference is successful. Hence in our example, the target (musical_piece1) is again replaced by a substitute metonym (music1). The source, which is music1, the object preference of play12, remains unchanged.

(Enter diamond 1 for a third time.) The sense-network path between the source (music1) and the latest target (music1) is sought. The path is inclusive, that music1 is a type of music1, so a literal relation is found.

(Exit the chart.) The processing of the preference-based semantic relation(s) between play12, and its preference for music1, and johann_sebastian_bach is completed.

After an initial preference violation (Johann Sebastian Bach is not a kind of music), the semantic relation found was an ARTIST FOR ART FORM metonymic relation (that johann_sebastian_bach composes musical pieces) followed by a CONTAINER FOR CONTENTS metonymic relation (that musical pieces contain music) followed by a literal relation (that music is music).

Example 20
"The car drank gasoline."

There is a metaphorical relation between car1 and the verb sense drink1 in (20). The source is drink1, whose agent preference is animal1, and the target is car1 (see Figure 6).

A metaphorical relation is sought after failing to find an inclusive network path or a metonymic inference between animal1 and car1, hence the network path between animal1 and car1 must be exclusive. The network path found is an estranged one.

The second stage is the match between the relevant cell of animal1 and the cells of car1. In the present example, drink1 is relevant. The list of cells for animal1 is searched for one referring to drinking. The relevant cell in the list is [animal1, drink1, drink1], which is matched against the inherited cells of car1 (see Figure 7). A sister match is found between [animal1, drink1, drink1] and [car1, use2, gasoline1] from car1.

The sister match is composed of two sister paths found in the sense-network. The first sister path is between the verb senses drink1 and use2, which are both types of expending (Figure 8). The second path is between the noun senses drink1 and gasoline1, which are both types of liquid (Figure 9). The effect of the network paths is to establish correspondences between the two cells such that an analogy is "discovered" that animals drink potable liquids as cars use gasoline. Note that, like Gentner's (1983) systematicity principle, the correspondences found are structural and independent of the content of the word senses they connect. Note also that the two cells have an underlying similarity or "ground" (Richards 1936) in that both refer to the expenditure of liquids. This second stage of finding a relevant analogy seems the crucial one in metaphor recognition.

Figure 10 shows the match of the nonrelevant cells from animal1 and car1. The cell [car1, use2, gasoline1] has been removed. There are three inclusive cell matches as animals and cars share physical objectlike properties of boundedness, three dimensions,

```
                sf(drink1,
                  [[arcs,
                    [[supertype, [ingest1, expend1]]]],
                   [node2,
                     [[agent,
                       [preference, animal1]],
                      [object,
                       [preference, drink1]]]]]).

sf(animal1,                              sf(car1,
   [[arcs,                                  [[arcs,
     [[supertype, organism1]]],               [[supertype, motor_vehicle1]]],
    [node0,                                  [node0,
     [[biology1, animal1],                     [[it1, carry1, passenger1]]]]).
      [it1, drink1, drink1],
      [it1, eat1, food1]]]]).
```

Figure 6
Sense-frames for drink1 (verb sense), animal1 and car1 (noun senses)

Relevant cell of animal1 (SOURCE)	Cells of car1 (TARGET)
[animal1, drink1, drink1]	[[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1], [animacy1, nonliving1], [car1, roll1, [on3, land1]], [composition1, steel1], [driver1, drive1, car1], [car1, have1, [4, wheel1]], [car1, have1, engine1], **[car1, use2, gasoline1],** [car1, carry1, passenger1]]

Figure 7
Match of relevant cell from animal1 against cells of car1

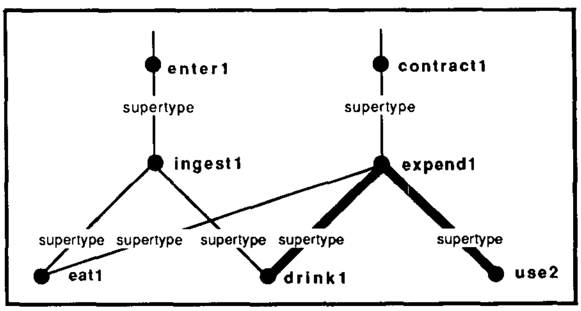

Figure 8
Sister sense-network path between drink1 and use2 (verb senses)

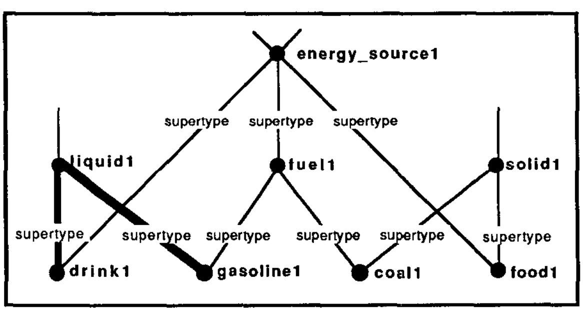

Figure 9
Sister sense-network path between gasoline1 and drink1 (noun senses)

Non-relevant cells of animal1 (SOURCE)	Non-relevant cells of car1 (TARGET)	Cell matches
[[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1],	[[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1],	3 same cell matches
[animacy1, living1], [composition1, flesh1],	[animacy1, nonliving1], [composition1, steel1],	2 sister cell matches
[animal1, eat1, food1], [biology1, animal1]]		2 distinctive source cells (of animal1)
	[car1, roll1, [on3, land1]], [driver1, drive1, car1], [car1, have1, [4, wheel1]], [car1, have1, engine1], [car1, carry1, passenger1]]	5 distinctive target cells (of car1)

Figure 10
Matches of non-relevant cells from animal1 and car1

[preference, [[network_path, [0, 0, 0, 0, 1]],	First array: preference violation (estranged sense-network path)
[cell_match, [[relevant, [0, 0, 1, 0, 0, 0, 10]],	Second array: relevant analogy (sister match of relevant cell)
[non_relevant, [0, 3, 2, 0, 0, 2, 5]]]]]]	Third array: distance betw. conceptual domains (matches of non-relevant cells)

Figure 11
Semantic vector for a metaphorical semantic relation

and solidity. Two cell matches are exclusive. Animals are composed of flesh, whereas cars are composed of steel. Animals are living, whereas cars are nonliving. There are two distinctive cells of animal1 and five distinctive cells of car1. Tourangeau and Sternberg's (1982) hypothesis predicts that the greater the distance between the conceptual domains of the terms involved in a metaphor, the more apt the metaphor. The proportion of similarities (inclusive cell matches) to differences (exclusive cell matches) is 3 to 2, which is a middling distance suggesting, tentatively, an unimposing metaphor.

All of these matches made by collation are recorded in the semantic vector shown in Figure 11. The crucial elements of the metaphorical relation in (20) are the preference violation and the relevant analogy. In Figure 11, the preference violation has been recorded as the 1 in the first array and the relevant analogy is the 1 in the second array. Information about the distance between conceptual domains is recorded in the third array.

The 'preference' label indicates that a preference has been matched (rather than an assertion). The five columns of the first array record the presence of ancestor, same, sister, descendant and estranged network paths respectively. When a preference is evaluated, only one network path is found, hence the single 1 in the fifth column, which indicates that an estranged network path was found between animal1 and car1.

Cell matches are recorded in the second and third arrays, which each contain seven columns. Those columns record the presence of ancestor, same, sister, descendant, estranged, distinctive source, and distinctive target cell matches respectively. The 1 in the third column of the second array is the relevant analogy — a sister match of the relevant cell [animal1, drink1, drink1] and the cell [car1, use2, gasoline1]. The 10 is the ten distinctive cells of car1 that did not match [animal1, drink1, drink1]. This is the match of 12 cells, 1 from the source and 11 from the target (see Figure 7). The sum of array columns is:

$$((0+0+1+0+0)\times 2)+((0+10)\times 1)=(1\times 2)+(10\times 1)=12.$$

The 3 similarities, 2 differences, 2 distinctive cells of animal1 and 5 distinctive cells of car1 are the nonzero numbers of the final array. The 3 similarities are all same cell matches; the 2 differences are both sister cell matches. A total of 17 cells are matched, 7 from the source and 10 from the target (see Figure 10). The total of array columns is:

$$((0+3+2+0+0)\times 2)+((2+5)\times 1)=(5\times 2)+(7\times 1)=17.$$

Example 15
"The ship ploughed the waves."

In (15), there is a metaphorical relation between a sense of the noun 'ship' and the second sense of the verb 'plough' in meta5's lexicon. Note that 'plough,' like 'drink,' belongs to several parts of speech. Figure 12 shows the sense-frames for the verb sense plough2, the noun sense plough1, which is the instrument preference of plough2, and the noun sense ship1.

In (15), meta5 matches senses of 'ship' against senses of 'plough.' When meta5 pairs ship1 with plough2, it calls upon collation to match ship1 against the noun sense plough1, the instrument preference of plough2.

First, the graph search algorithm searches the sense-network for a path between plough1 (which is the preference) and ship1 and finds an estranged network path between them, i.e., a ship is not a kind of plough, so plough2's instrument preference is violated.

Next, collation inherits down lists of cells for plough1 and ship1 from their superordinates in the sense-network. What is relevant in the present context is the action of ploughing because (15) is about a ship *ploughing* waves. Collation then runs through the list of inherited cells for the noun sense plough1 searching for a cell that refers to the action of ploughing in the sense currently under examination by meta5, plough2.

```
sf(plough2,
   [[arcs
     [[supertype, transfer1]]],
    [node2,
     [[instrument,
       [preference, plough1]],
      [object,
       [preference, soil1]]]]]).

sf(plough1,                        sf(ship1,
   [[arcs,                            [[arcs,
     [[supertype, tool1]]],             [[supertype, watercraft1]]],
    [node0,                            [node0,
     [[farmer1, plough1, it1],          [[it1, carry1, shipment1],
      [it1, plough2, soil1]]]]]).        [it1, have1, engine1]]]]]).
```

Figure 12
Sense-frames for plough2 (verb sense), plough1 and ship1 (noun senses)

Relevant cell of plough1 (SOURCE)	Cells of ship1 (TARGET)
[plough1, plough2, soil1]	[[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1], [animacy1, nonliving1], [composition1, metal1], [ship1, use2, energy_source1], [boatman1, sail1, ship1], **[ship1, sail2, water2],** [ship1, carry1, shipment1], [ship1, have1, engine1]]

Figure 13
Match of relevant cell from plough1 against cells from ship1

Non-relevant cells of ploughl (SOURCE)	Non-relevant cells of ship1 (TARGET)	Cell matches
[[composition1, matter1],	[[composition1, metal1],	1 ancestor cell match
[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1], [animacy1, nonliving1],	[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1], [animacy1, nonliving1],	4 same cell matches
[farmer1, plough1, plough1]]	[boatman1, sail1, ship1],	1 sister cell match
	[ship1, use2, energy_source1], [ship1, carry1, shipment1], [ship1, have1, engine1]]	3 distinctive target cells (of ship1)

Figure 14
Matches of non-relevant cells from plough1 and ship1

Collation finds a relevant cell [plough1, plough2, soil1] and uses its frame-matching algorithm to seek a match for the cell against the list of inherited cells for ship1, shown in Figure 13 (for ease of reading, it1 has again been replaced by the word senses being defined). The algorithm finds a match with [ship1, sail2, water2] (highlighted in Figure 13), and hence collation "discovers" a relevant analogy that both ships and ploughs move through a medium, i.e., that ploughs plough through soil as ships sail through water.

Finally, collation employs the frame matching algorithm a second time to match together the remaining nonrelevant cells of plough1 and ship1 (see Figure 14). The cell [ship1, sail2, water2] is removed to prevent it from being used a second time.

Figure 15 shows the semantic vector produced. As with Figure 11, it shows a metaphorical relation. There is a preference violation, an estranged network path indicated by the 1 in the fifth column of the first array. There is also a relevant analogy, shown by the 1 in the third column of the second array: the analogical match of the cells [plough1, plough2, soil1] and [ship1, sail2, water2]. The second array shows that 11 cells are matched, 1 from the source and 10 from the target (check against Figure 13). The sum of the array's columns is:

$$((0+0+1+0+0)\times 2)+((0+9)\times 1)=(1\times 2)+(9\times 1)=11.$$

```
[preference,
  [[network_path,
    [0, 0, 0, 0, 1]],
   [cell_match,
    [[relevant,
      [0, 0, 1, 0, 0, 0, 9]],
     [non_relevant,
      [1, 4, 1, 0, 0, 0, 3]]]]]]
```

`[[network_path, [0, 0, 0, 0, 1]]`	First array: preference violation (estranged sense-network path)
`[relevant, [0, 0, 1, 0, 0, 0, 9]]`	Second array: relevant analogy (sister match of relevant cell)
`[non_relevant, [1, 4, 1, 0, 0, 0, 3]]`	Third array: distance betw. conceptual domains (matches of non-relevant cells)

Figure 15
Semantic vector for another metaphorical semantic relation

In the third array, the match of nonrelevant cells, there is 1 ancestor match, 4 same matches, 1 sister match, and 3 distinctive cells of ship1. Fifteen cells are matched, 6 from the source and 9 from the target (see Figure 14). The totals are:

$$((1 + 4 + 1 + 0 + 0) \times 2) + ((0 + 3) \times 1) = (6 \times 2) + (3 \times 1) = 15.$$

Semantic vectors can represent all the semantic relations except metonymic ones. The reason is that metonymic relations, unlike the others, are not discriminated by CS in terms of only five kinds of network path and seven kinds of cell matches. Instead, they consist of combinations of network paths and specialized matches of cells that have not fallen into a regular enough pattern to be represented systematically.

6. Extensions

Even for those semantic dependencies investigated, the interpretation of semantic relations seems to require more complexity than has been described so far in this paper. Consider the differences between the following sentences:

Example 20
"The car drank gasoline."

Example 27
"The car drank coffee."

Intuitively, sentence (20) is metaphorical while (27) is metaphorical/anomalous. In (20), the semantic relation between 'car' and 'drink' is thought to be metaphorical, and the isolated semantic relation between just 'drink' and 'gasoline' is anomalous, but the sentence as a whole is metaphorical because it is metaphorical that cars should use up gasoline.

In (27), the semantic relation between 'car' and 'drink' is metaphorical; the semantic relation between just 'drink' and 'coffee' is literal; yet the effect of (27) as a whole is metaphorical/anomalous. The object preference of 'drink' is for a drink, i.e., a potable liquid. It seems that it is metaphorical for cars to "drink" a liquid commonly used up by cars, e.g., gasoline, but anomalous if the liquid has nothing to do with cars, e.g., coffee, as in (27).

The problem of understanding the differences between sentences (20) and (27) requires some further observations about the nature of semantic relations, principally that the differences are caused by the combinations of semantic relations found in the sentences and the *relationships* between those relations. Below is a suggestion as to how deeper semantic processing might discriminate the differences between the two sentences.

Before getting to the deeper processing, we need a better semantic vector notation. The better semantic vector notation, which developed from a discussion with Afzal Ballim, is a modification of the notation shown in Section 5. The key differences are *reformulation* by rewriting the five and seven column arrays in terms of the predicate-argument notation used in the rest of semantic vectors, and *extension* by adding the domain knowledge connected by every network path and cell match.

Figure 16 shows the semantic vector in Figure 11 reformulated and extended. The advantage of vectors like the one in Figure 16 is that they record both how the sense-frames of two word senses are matched (i.e., as various kinds of network path and cell match) and what information in the sense-frames is matched (i.e., all the cells). For example, the part of Figure 16 that begins "[relevant, ..." contains all the information found in Figure 7, the match of the relevant cell from animal1 against the cells of car1, both the types of cell matches *and the cells matched.* The equivalent part of Figure 11 only records the types of cell matches. Recording the contents of the matched cells is useful because it enables a deepened analysis of semantic relations. Such an analysis is needed to detect the differences between (20) and (27).

In the description of CS in Section 4, collation discriminates the one or more semantic relations in each semantic dependency, but treats the semantic relations in one dependency as isolated from and unaffected by the semantic relations in another dependency. What is needed is extra processing that interprets the semantic relation(s) in a later dependency with respect to the semantic relation(s) established in an earlier

```
[preference,
 [[network_path,
     [estranged,
         [1, [animal1, car1]]]],
  [cell_match,
   [[relevant,
     [[sister,
         [1, [[[animal1, drink1, drink1], [car1, use2, gasoline1]]]],
      [distinctive_target,
         [10, [[bounds1, distinct1], [extent1, three_dimensional1],
               [behaviour1, solid1], [composition1, metal1],
               [animacy1, nonliving1], [car1, roll1, [on3, land1]],
               [driver1, drive1, car1], [car1, have1, [4, wheel1]],
               [car1, have1, engine1], [car1, carry1, passenger1]]]]]],
    [non_relevant,
      [[same,
         [3, [[[bounds1, distinct1], [bounds1, distinct1]],
              [[extent1, three_dimensional1], [extent1, three_dimensional1]],
              [[behaviour1, solid1], [behaviour1, solid1]]]]],
       [sister,
         [2, [[[composition1, flesh1], [composition1, metal1]],
              [[animacy1, living1], [animacy1, nonliving1]]]]],
       [distinctive_source,
         [2, [[animal1, eat1, food1], [biology1, animal1]]]],
       [distinctive_target,
         [5, [[car1, roll1, [on3, land1]], [driver1, drive1, car1],
              [car1, have1, [4, wheel1]], [car1, have1, engine1],
              [car1, carry1, passenger1]]]]]]]]]]
```

Figure 16
Reformulated and extended version of Figure 11

```
[preference,
  [cell_match,
    [relevant,
      [sister,
        [1, [[animal1, drink1, drink1], [car1, use2, gasoline1]]]]]]
```

Figure 17
Vector statement of match of relevant cell from animal1 against cells of car1

```
[preference,
  [cell_match,
    [relevant,
      [sister,
        [1, [[animal1, drink1, drink1], [vehicle1, use2, gasoline1]]]]]]
```

Figure 18
Vector statement of match of relevant cell from drink1 against cells of gasoline1 (noun senses)

one. This processing matches the domain knowledge in semantic vectors, i.e., this processing is a comparison of coherence representations.

In sentences such as (20) and (27) there are two key semantic dependencies. The first one is between the subject noun and the verb; the second is between the verb and object noun. In each dependency, the source is the verb (through its agent and object preferences) and the targets are the nouns. Semantic relations are found for each dependency. One way to detect the difference between metaphorical sentences such as (20) and metaphorical/anomalous ones such as (27) is in each sentence to consult the semantic vectors produced in its two main semantic dependencies and compare the matches of the relevant cells that are found by collation.

Let us go through such an analysis using CS, starting with the first semantic dependency between subject noun and verb. In this semantic dependency in both (20) and (27), a relevant analogy is discovered as part of a metaphorical relation between the target car1 and animal1, the agent preference of the source drink1. The semantic vector in Figure 16 records the two cells that figure in that relevant analogy. Figure 17 shows the same information from the semantic vector but written as a statement.

When the second semantic dependency is analyzed in (20), the target is gasoline1 and is matched against the noun sense drink1, the object preference of the source drink1 (the verb sense). A semantic vector is produced. The relevant cell found in the noun sense drink1 is [animal1, drink1, drink1]. Its match against [vehicle1, use2, gasoline1], a cell from gasoline1, is shown in the vector statement in Figure 18. The match is a sister match, indicating a relevant analogy.

Now this is peculiar because "drinking gasoline" is anomalous, yet a relevant analogy has been found and this paper has argued that relevant analogies are special to metaphorical relations. One possible explanation is that differences exist between the recognition of metaphorical relations that concern agents and metaphorical relations that concern objects and other case roles. It may be that metaphorical relations are indicated by a relevant analogy, but only *in selected circumstances*. This needs further investigation.

```
[preference,
  [cell_match,
    [relevant,
      [ancestor,
        [1, [[animal1, drink1, drink1], [human_being1, drink1,  coffee1]]]]]]]
```

Figure 19
Vector statement of match of relevant cell from drink1 against cells from coffee1 (noun senses)

To return to the analysis of (20), what appears to be important in determining that (20) is a metaphorical sentence is the comparison of the two pairs of matched relevant cells:

[[animal1, drink1, drink1], [car1, use2, gasoline1]]

[[animal1, drink1, drink1], [vehicle1, use2, gasoline1]]

The two source cells are the same and the two target cells, [car1, use2, gasoline1] and [vehicle1, use2, gasoline1], are almost identical, indicating that the same basic analogy runs through the whole of (20), hence the sentence as a whole is metaphorical.

Now let us analyze the second semantic dependency in (27). The target is coffee1 and is again matched against drink1, the object preference of the verb sense drink1, the source. The relevant cell from the noun sense drink1 is again [animal1, drink1, drink1], which matches against [human_being1, drink1, coffee1] from the target coffee1. This time, the match is an ancestor match and hence not a relevant analogy. Figure 19 shows this match of the relevant cell as a vector statement. Let us compare the two pairs of matched relevant cells for (27):

[[animal1, drink1, drink1], [car1, use2, gasoline1]]

[[animal1, drink1, drink1], [human_being1, drink1, coffee1]]

The two source cells are the same but the two target cells, [car1, use2, gasoline1] and [human_being1, drink1, coffee1], are very different. The reason that the sentence as a whole is metaphorical/anomalous is because of the clash between these target cells. The basic analogy of a car ingesting a liquid does not carry over from the first semantic dependency into the second. The anomalous flavor of (27) could not be detected by looking at the semantic relations in the dependencies in isolation because one semantic relation is metaphorical and the other is literal. Neither relation is anomalous — the anomaly comes from the interaction between the two relations.

Figure 20 is a proposed representation for sentence (20). The left side of Figure 20 shows the knowledge representation part of the sentence representation: a simple case-frame based representation of (20). The right side of Figure 20, within the grey partition, is the coherence representation component of the sentence representation: abridged semantic vectors for the two main semantic dependencies in (20). The upper semantic vector is the match of the target car1 against the source animal1. The lower semantic vector is the match of the target gasoline1 against the source drink1, the noun sense. The upper abridged semantic vector indicates a metaphorical relation. The lower semantic vector also indicates a metaphorical relation though, as was noted earlier, "drinking gasoline" when interpreted in isolation is surely anomalous.

The underlines in Figure 20 denote pointers linking the semantic vectors to the case frame. The grey vertical arrows show that the two semantic vectors are also linked

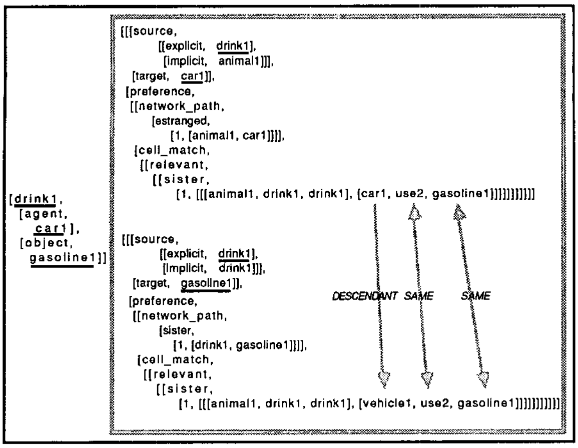

Figure 20
Sentence representation for "The car drank gasoline"

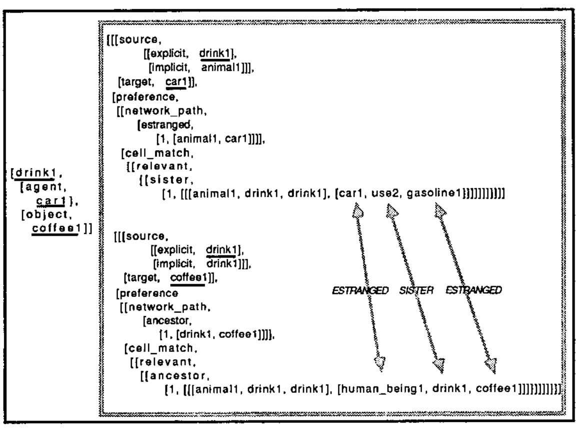

Figure 21
Sentence representation for "The car drank coffee"

together via the matches of their relevant cells. In those matches, the arrows are sense-network paths found between the elements of the two target cells. The network paths indicated in grey, that connect the two abridged semantic vectors, show *processing of coherence representations*. The particular network paths found (indicated in italics), a descendant path and two same "paths," show that the same relevant analogy is used in both semantic relations — that both semantic relations involve a match between animals drinking potable liquids and vehicles (including cars) using gasoline — hence sentence (20) as a whole is metaphorical. Figure 20 is therefore unlike any of the coherence representations shown previously, because it shows a representation of a metaphorical *sentence*, not just two isolated metaphorical relations.

Compare Figure 20 with Figure 21, a sentence representation for (27). The upper semantic vector again indicates a metaphorical relation between car1 and drink1. The lower semantic vector indicates a literal relation between drink1 and coffee1. What is important here is the match of relevant information discovered in the two semantic relations, as indicated by the three network paths. The paths found are two estranged paths and a sister path, indicating that the relevant information found during the two semantic relations is different: in one semantic relation, information about animals drinking potable liquids is matched against *cars using gasoline*; in the other, the same information is matched against *human beings drinking coffee*; but cars using gasoline and human beings drinking coffee are quite different, hence sentence (27) is anomalous overall.

Note that in Figures 20 and 21, the coherence representation part of the sentence representation is much larger than the knowledge representation part. The detailed "world knowledge" about car1, the verb sense drink1, gasoline1, and coffee1 are all on the right side. It is interesting to contrast the figures with early Conceptual Dependency (CD) diagrams such as those in Schank (1973) because, rather than the large and seemingly unlimited amounts of world knowledge that appear in CD diagrams, the two figures present only the world knowledge needed to discriminate the semantic relations in (20) and (27).

7. Discussion and Conclusions

This section reviews the material on metonymy and metaphor in Section 2 in light of the explanation of the met* method given in Sections 3–6. When compared with the AI work described in Section 2, the met* method has three main advantages. First, it contains a detailed treatment of metonymy. Second, it shows the interrelationship between metonymy, metaphor, literalness, and anomaly. Third, it has been programmed.

Preference Semantics addresses the recognition of literal, metaphorical, and anomalous relations, but does not have a treatment of metonymy. In the case of Preference Semantics, the theory described in Wilks (1978) has not been implemented, though the projection algorithm was implemented (Modiano 1986) using some parts of CS to supply detail missing from Wilks' original specification.

Gentner's (1983) Structure-Mapping Theory has no treatment of metonymy. The theory has been implemented in the Structure-Mapping Engine (Falkenhainer, Forbus and Gentner 1989) and some examples analyzed by it but not, to my knowledge, examples of metaphor or anomaly.

Indurkhya's (1988) Constrained Semantic Transference theory of metaphor has no treatment of metonymy, anomaly or literalness. It has also not been implemented: see Indurkhya (1987) for reasons why.

Hobbs and Martin (1987) offer a relatively shallow treatment of metonymy without, for instance, acknowledgement that metonymies can be driven from either the source or the target. Hobbs' "selective inferencing" approach to text interpretation has been applied to problems including lexical ambiguity (Hobbs 1977; 1982b; Hobbs and Martin 1987), metaphor (Hobbs 1977; 1983a; 1983b) and the "local pragmatics" phenomena of metonymy (Hobbs and Martin 1987), but not anomaly. To my knowledge, Hobbs has yet to produce a unified description of selective inferencing that shows *in detail* how lexical ambiguity is resolved or how the differences between metaphor, metonymy, and so on can be recognized. Hobbs' earlier papers include a series of programs — SATE, DIANA, and DIANA-2 — but the papers are not clear about what the programs can do. It is not clear, for example, whether any of the programs actually analyze any metaphors.

Martin's (1990) work is the only other computational approach to metaphor that has been implemented. However, the work does not have a treatment of metonymy. Martin's metaphor-maps, which are used to represent conventional metaphors and the conceptual information they contain, seem to complement semantic vectors of the extended kind described in Section 6. In Section 6, I argued that vectors need to record the conceptual information involved when finding mappings between a source and target. What metaphor-maps do is freeze (some of) the conceptual information involved in particular metaphorical relations. There is some theoretical convergence here between our approaches; it would be interesting to explore this further.

Moreover, the metaphors studied so far in CS seem linked to certain conventional metaphors because certain types of ground have recurred, types which resemble Lakoff and Johnson's (1980) structural metaphors. Two types of ground have cropped up so far.

Example 28
"Time flies."

The first is a use-up-a-resource metaphor which occurs in (20) and in (28) when viewed as noun-verb sentence. Both sentences are analyzed by meta5. Use-up-a-resource resembles structural metaphors like TIME IS A RESOURCE and LABOR IS A RESOURCE which, according to Lakoff and Johnson (1980, p. 66), both employ the simple ontological metaphors of TIME IS A SUBSTANCE and AN ACTIVITY IS A SUBSTANCE:

> These two substance metaphors permit labor and time to be quantified — that is, measured, conceived of as being progressively "used up," and assigned monetary values; they allow us to view time and labor as things that can be "used" for various ends.

Example 29
"The horse flew."

The second type of ground is motion-through-a-medium, a type of ground discussed by Russell (1976). This appears in (15) and (29), again both analyzed by meta5.

Incidentally, it is worth noting that structural metaphors have proven more amenable to the met* method than other kinds tried. I assumed initially that orientational and ontological metaphors would be easier to analyze than structural metaphors because they were less complex. However, structural metaphors have proved easier to analyze, probably because structural metaphors contain more specific concepts such as "drink" and "plough," which are more simple to represent in a network structure (like the sense-network of CS) so that analogies can be found between those concepts.

7.1 Relationship between Literalness and Nonliteralness

We return here to Gibbs' point concerning the traditional notion of literal meaning that [1] all sentences have literal meanings that are entirely determined by the meanings of their component words and that [2] the literal meaning of a sentence is its meaning independent of context. Although [1] and [2] are both presently true of CS, there are means by which context can be introduced more actively into sentence interpretation.

At present, the meaning of a sentence in CS — whether literal or nonliteral — is not derived entirely independently of context; however, the only context used is a limited notion of relevance which is generated by collation from within the sentence being analyzed: what is relevant is given by the sense of the main sentence verb. Nevertheless, because of this notion of relevance, contextual influence *is* present in semantic interpretation in CS. Moreover, the notion of relevance is recorded in semantic vectors (Figures 11 and 15) and the extended coherence representations discussed in Section 6. Hence, the processes and representations of CS possess basic equipment for handling further kinds of context.

7.2 Relationship between Metonymy and Metaphor

The met* method is consistent with the view that metaphor is based on similarity, whereas metonymy is based on contiguity (cf. Jakobsen and Halle 1956). Contiguity, readers may recall, refers to being connected or touching whereas similarity refers to being alike in essentials or having characteristics in common. The difference comes from *what* and *how* the conceptual information is related.

Example 1
"My car drinks gasoline."

Let us consider *what* is related first. In metaphor, an *aspect* of one concept is similar to an *aspect* of another concept; e.g., in (1), an aspect of the concept for animal, that animals drink potable liquids, is similar to an aspect of another concept, that cars use gasoline.

Example 2
"The ham sandwich is waiting for his check."

However, in metonymy, a whole *concept* is related to an *aspect* of another concept. For example, in (2) the metonymy is that the concept for ham sandwich is related to an aspect of another concept, for "the man who ate a ham sandwich."

Regarding *how* that conceptual information is related: in the case of metaphor, the met* method assigns a central role to finding an analogy, and an analogy between two terms is due to some underlying similarity between them (the ground), e.g., in the analogy that animals drinking potable liquids is like cars using gasoline, the underlying similarity is that both animals and cars ingest liquids. In an analogy, the relationship between aspects of two concepts is purely structural. In metonymies, however, the relationships are "knowledge-laden" connections, e.g., PART-WHOLE and CONTAINER-CONTENTS.

So in summary, "similarity" in metaphor is understood to be based on structural relationships between aspects of concepts, whereas "contiguity" in metonymy is based on knowledge-specific relationships between a concept and an aspect of another concept. These observations, I would argue, support the view that metonymy has primarily a referential function, allowing something to *stand* for something else — a connection between a concept and an aspect of another concept. The observations also support the view that metaphor's primary function is understanding, allowing something to be conceived of in terms of something else: the role of analogy is especially crucial to this function.

7.3 Metonymy

The treatment of metonymy permits chains of metonymies (Reddy 1979), and allows metonymies to co-occur with instances of either literalness, metaphor, or anomaly.

The kinds of inferences sought resemble the kinds of inferences that Yamanashi (1987) notes link sentences. An obvious direction in which to extend the present work is toward across-sentence inferences.

Example 30
"John drank from the faucet" (Lehnert 1978, p. 221).

Example 31
"John filled his canteen at the spring" (Ibid.).

Metonymy seems closely related to the work on non-logical inferencing done by Schank (Schank 1973) and the Yale Group (Schank 1975; Schank and Abelson 1977; Schank and Riesbeck 1981). For example, Lehnert (1978) observes that just one inference is required for understanding both (30) and (31). The inference, that water comes from the faucet in (30) and the spring in (31), is an instance of PRODUCER FOR PRODUCT in which the faucet and spring are PRODUCERs and water is the PRODUCT. However, the inference is not a metonymy because it is from unused cases of the verbs 'drink' and 'fill' whereas metonymy only occurs in the presence of a violated selection restriction, that neither (30) nor (31) contain.

7.4 Metaphor

Metaphor recognition in the met* method is related to all four views of metaphor described in Section 2, consisting of:

1. a contextual constraint violation, such as a preference violation — as in the selection restrictions view;
2. a set of "correspondences" — rather like the system of commonplaces in the interaction view;
3. a relevant analogy — cf. the comparison and interaction views; with
4. analogies that fall into patterns not unlike conceptual metaphors found in the conventional view.

In CS, the presence of metaphor has been investigated in violations of preferences, a kind of lexical contextual constraint. Though clearly this is a small part of the picture, it seems worth establishing an extensive picture of preference violation and metaphor before moving on to other contextual constraints.

Collation and the met* method have certain similarities with the comparison view of metaphor, especially in the cell matching process. The relevant analogies discovered in CS are indeed, to quote Tourangeau and Sternberg, "a comparison in which one term... is asserted to bear a partial resemblance to something else."

The collation process gives quite a clear picture of the ground and tension in a metaphor. The ground is the most specific statement that subsumes both statements that figure in the analogy, e.g., [it1, ingest1, liquid1] is the ground for the analogy involving [animal1, drink1, drink1] and [car1, use2, gasoline1] (see Figures 8 and 9). Moreover, the details of the process match well Aristotle's two basic principles for finding the ground of a metaphor in that both terms in a metaphorical relation belong to a common category (in the example above, the common categories are it1, ingest1, and liquid1) *and* an analogy is found between them.

The collation process also takes care of many of the problems Tourangeau and Sternberg (1980) note with the comparison view. Regarding the problem that "everything shares some feature or category... with everything else," CS is in agreement: the only significant combination of features in a metaphor are those involved in a relevant analogy. The problem that "the most obvious shared features are often irrelevant," i.e., that the most obvious shared features are irrelevant to a metaphor, is borne out by experience with CS — for example, animals and cars share some basic physical object-like properties, but these have a minor role in understanding cars drinking. The met* method bears out another problem that, "even when a feature is relevant, it is often shared only metaphorically." Finally, with the problem that novel metaphors cannot be based on "extant similarities," — the relevant analogies found in the met* method are not "extant" but have to be actively discovered.

In Section 2, two main differences were noted between the interaction and comparison views: first, that similarities are "created" in the interaction view, whereas only pre-existing similarities are found in the comparison view, and second, that a whole system of similarities are evoked in the interactions view, unlike the comparisons view, which focuses upon finding a single similarity. Regarding the first difference, I would argue that the difference is a mistaken one and that interaction theorists are simply using a sophisticated form of comparison. This is quite evident when one examines, for example, the methods Tourangeau and Sternberg propose for relating features across domains in their theory. The second of Aristotle's basic principles is finding an analogy, yet Tourangeau and Sternberg (1982, p. 218) themselves say that, "in a sense, we are proposing that metaphors are analogies that include both tenor and vehicle and their different domains as terms."

And, of course, finding an analogy is central to the met* method on CS.

Regarding the second difference, I would agree that finding a system of commonplaces is distinctive. However, the extensions to CS described in Section 6 move toward the direction of finding a system of commonplaces in that the deeper semantic vectors, and sentence representations shown in Figures 20 and 21 contain the information crucial to finding a system of commonplaces. Having identified the crucial analogy in (20), the deeper semantic vector contains the two pairs of matched relevant cells that provide the core analogy on which the metaphorical interpretation of (20) is built:

[[animal1, drink1, drink1], [car1, use2, gasoline1]]

[[animal1, drink1, drink1], [vehicle1, use2, gasoline1]]

With this information at hand, the sense-frames for word senses in analogical correspondence — the verb senses drink1 and use2, the noun senses animal1 and car1, animal1 and vehicle1, and drink1 and gasoline1 — can be systematically expanded to uncover deeper commonplaces between animals and cars.

In conclusion, the view of metonymy and metaphor in the met* method is consistent with much of the literature on these phenomena. The met* method is consistent with the view that the primary function of metaphor is understanding while that of metonymy is referential, like anaphora. Nevertheless, metonymy and metaphor do have much in common: both might be described as forms of "conceptual ellipsis," a shorthand way of expressing ideas.

The met* method in its present serial form recognizes literalness, metonymy, metaphor, and anomaly in the following order and by the following characteristics.

- Literalness — a satisfied preference.
- Metonymy — a successful conceptual inference.
- Metaphor — an underlying relevant analogy.
- Anomaly — none of the above.

The above analysis also illustrates, I hope, why metonymy and metaphor are easily confused: both are nonliteral and are found through the discovery of some aspect (a property) shared by the source, a preference, and the target, in the above case a surface noun. The differences are (a) how that aspect is selected, (b) the operations that follow, (c) the effect those operations produce, and (d) subsequent processing.

In the case of metonymy, (a) the selected aspect forms a regular semantic relationship with a property from the target; (b) there is substitution, i.e., replacement of one concept with another; (c) hence the apparent referential function of metonymy; and (d) is unclear at present.

In the case of metaphor, (a) the selected aspect is relevant; (b) forms an analogy with another aspect from the target; and (c) the effect is of surprise discovery of similarity between the two concepts; and (d) the discovered analogy is used to unearth further similarities between the two concepts (i.e, to deepen the analogy) and to guide subsequent sentence interpretation. Moreover, the view of metaphor in CS contains elements of the selection restrictions view, the comparisons view, and the interactions view of metaphor.

It should be emphasized that the met* method has only been applied to a small set of English sentences. Metonymy interpretation has been investigated only for adjective-noun and subject-verb-object constructions; metaphor interpretation, only for the latter. The best avenue for progress with the met* method appears to be the extensions to metaphor interpretation described in Section 6. In the meantime I am looking for sentences that contain semantic relations consisting of a metonymy (or chain of metonymies) followed by a metaphor.

Example 32
"America believes in democracy" (Hobbs 1983b, p. 134).

On a related point, some sentences are interesting in this respect because they have either a metaphorical or metonymic interpretation. In (32), for example, "Are we viewing America metaphorically as something which can believe, or are we using it metonymically to refer to the typical inhabitant, or the majority of inhabitants, of America?" (Ibid., p. 135).

Example 33
"Prussia invaded France in 1870."

Sentence (33), which was discussed in a group working on beliefs at the CRL (see Acknowledgments), also has separate metonymic and metaphorical interpretations. The key semantic relation is between 'Prussia' and 'invade.' The relation is nonliteral because 'army' is the expected agent of 'invade' and 'Prussia' is a country, not an army. What, then, is the semantic relation between 'Prussia' and 'army'? One possibility is that a chain of metonymies is involved, that the army is controlled by the government which also controls Prussia. A second possibility is that Prussia is understood metaphorically as being an animate thing that extends itself into France.

Acknowledgments

I would like to thank the many people at the Cognitive Studies Centre, University of Essex; the Computing Research Laboratory, New Mexico State University; and the Centre for Systems Science, Simon Fraser University, with whom I have had fruitful discussions over the years, especially those in the beliefs group at the CRL (Afzal Ballim, John Barnden, Sylvia Candelaria de Ram, and Yorick Wilks); others at the CRL (including Xiuming Huang, David Farwell, and Eric Dietrich); and colleagues in the CSS who made helpful comments on earlier drafts of this paper (Chris Groeneboer, Gary Hall, and Carl Vogel). A special word of thanks for the help given by Yorick Wilks, the director of the CRL, and Nick Cercone, the director of the CSS. I also gratefully acknowledge the financial support provided by SERC Project GR/C/68828 while at Essex, by the New Mexico State Legislature while at NMSU, and by the Advanced Systems Institute and the Centre for Systems Science while at SFU.

References

Amsler, Robert A. (1980). *The structure of the Merriam-Webster Pocket Dictionary.* Doctoral dissertation TR-164, University of Texas, Austin, Texas.

Basso, K.H. (1976). " 'Wise words' of the Western Apache: metaphor and semantic theory." In *Meaning in anthropology,* edited by K.H. Basso and H. Selby, New Mexico Press.

Beardsley, M. (1962). "The metaphorical twist." *Philosophy and Phenomenological Research,* 22:293–307.

Bickerton, Derek (1969). "Prolegomena to a linguistic theory of metaphor." *Foundations of Language,* 5:36–51.

Black, Max (1979). "More about metaphor." In *Metaphor and thought,* edited by A. Ortony, 19–43. Cambridge University Press.

Black, Max (1962). *Models and Metaphors.* Cornell University Press.

Boguraev, Branimir K. (1979). "Automatic resolution of linguistic ambiguities." Technical Report No. 11, University of Cambridge Computer Laboratory, Cambridge, England.

Campbell, P. (1975). "Metaphor and linguistic theory." *Quarterly Journal of Speech,* 61:1–12.

Chomsky, Noam (1964). "Degrees of grammaticalness." In *The Structure of Language: Readings in the Philosophy of Language,* edited by J.A. Fodor and J.J. Katz, 384–389, Prentice-Hall.

Van Dijk, T. (1975). "Formal semantics of metaphorical discourse." *Poetics,* 4:173–198.

Dubois, J.; Edeline, F.; Klinkenberg, J.-M.; Minguet, P.; Pire, F.; and Trinon, H. (1970). *Rhétorique Générale.* Larousse. Published as: *A General Rhetoric,* translated by P.B. Burrell and E.M. Slotkin, 1981, The John Hopkins University Press.

Falkenhainer, Brian; Forbus, Kenneth D.; and Gentner, Dedre (1986). "The structure-mapping engine: algorithm and examples." *Artificial Intelligence,* 41:1–63.

Fass, Dan C. (1986). "Collative semantics: an approach to coherence." Technical Report MCCS-86-56, Computing Research Laboratory, June 1986.

Fass, Dan C. (1987). "Semantic relations, metonymy, and lexical ambiguity resolution: a coherence-based account." In *Proceedings, 9th Annual Cognitive Science Society Conference,* University of Washington, Seattle, Washington: 575–586.

Fass, Dan C. (1988a). *Collative semantics: a semantics for natural language processing.* Doctoral dissertation, New Mexico State University, New Mexico.

Fass, Dan C. (1988b). "An account of coherence, semantic relations, metonymy, and lexical ambiguity resolution." In *Lexical Ambiguity Resolution: Perspectives from Psycholinguistics, Neuropsychology, and Artificial Intelligence,* edited by Steven L. Small, Garrison W. Cottrell, and Michael K. Tanenhaus, Morgan Kaufmann, 151–178.

Fass, Dan C. (1988c). "Metonymy and metaphor: what's the difference?" In *Proceedings, 12th International Conference on Computational Linguistics (COLING-88),* Budapest, Hungary, 177–181.

Fass, Dan C. (1988d). "Collative semantics: a study in the discrimination of meaning." Technical Report CSS/LCCR TR 88-24, Centre for Systems Science, Simon Fraser University.

Fass, Dan C. (1989a). "Lexical semantic constraints." Technical Report CSS/LCCR TR 89-11, Centre for Systems Science, Simon Fraser University. (To appear in

Semantics and the Lexicon, edited by James Pustejovsky. Kluwer Academic Publishers.)

Fass, Dan C. (1989b). "Four general representations and processes for use in problem solving." In *Knowledge Based Computer Systems (Proceedings of KBCS '89, Bombay, India)*, edited by S. Ramani, R. Chandrasekar, and K.S.R. Anjeyulu, Narosa Publishing House, 169–178.

Genette, Gerard (1970). "La rhetorique restreinte." *Communications*, 16:158–171.

Gentner, Dedre (1983). "Structure mapping: a theoretical framework for analogy." *Cognitive Science*, 7:155–170.

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

Grosz, Barbara J.; Appelt, Douglas E.; Martin, Paul; and Pereira, Fernando C.N. (1987). "TEAM: An experiment in the design of transportable natural-language interfaces." *Artificial Intelligence*, 32(2):173–243.

Harris, R. (1976). "Comprehension of metaphor: a test of a two-stage processing model." *Bulletin of the Psychonomic Society*, 8:321–324.

Hesse, M. (1966). *Models and Analogies in Science*. Notre Dame University Press.

Hobbs, Jerry R. (1977). "Coherence and interpretation in English texts." In *Proceedings, 5th International Joint Conference on Artificial Intelligence (IJCAI-77)*, Cambridge, Massachusetts, 110–116.

Hobbs, Jerry R. (1979). "Coherence and coreference." *Cognitive Science*, 3:67–90.

Hobbs, Jerry R. (1980). "Selective inferencing." In *Proceedings, 3rd National Conference of the Canadian Society for Computational Studies of Intelligence (CSCSI-3)*, Victoria, British Columbia, Canada, 101–114.

Hobbs, Jerry R. (1982a). "Towards an understanding of coherence in discourse." In *Strategies for Natural Language Processing*, edited by Wendy G. Lehnert and Martin H. Ringle, Erlbaum Associates, 223–243.

Hobbs, Jerry R. (1982b). "Representing ambiguity." In *Proceedings, 1st West Coast Conference on Formal Linguistics*, Stanford, California, 15–28.

Hobbs, Jerry R. (1983a). "Metaphor interpretation as selective inferencing: cognitive processes in understanding metaphor (Part 1)." *Empirical Studies of the Arts*, 1:17–33.

Hobbs, Jerry R. (1983b). "Metaphor interpretation as selective inferencing: cognitive processes in understanding metaphor (Part 2)." *Empirical Studies of the Arts*, 1:125–141.

Hobbs, Jerry R. (1985). "On the coherence and structure of discourse." Technical Report No. CSLI-85-37, Center for the Study of Language and Information (CSLI), Stanford University.

Hobbs, Jerry R. and Martin, Paul (1987). "Local Pragmatics." In *Proceedings, 10th International Joint Conference on Artificial Intelligence (IJCAI-87)*, Milan, Italy, 520–523.

Huang, Xiuming (1985). "Machine translation in the SDCG (Semantic Definite Clause Grammars) formalism." In *Proceedings, Conference on Theoretical and Methodological Issues in Machine Translation of Natural Languages*, Colgate University, New York, New York, 135–144.

Indurkhya, Bipin (1985). "A computational theory of metaphor comprehension and analogical reasoning." Technical Report #85/001, Boston University.

Indurkhya, Bipin (1987). "Approximate semantic transference: a computational theory of metaphors and analogies." *Cognitive Science*, 11(4):445–480.

Indurkhya, Bipin (1988). "Constrained semantic transference: a formal theory of metaphors." In *Analogica: Proceedings of the First Workshop on Analogical Reasoning*, edited by Armand Prieditis, Morgan Kaufmann (Pitman Research Notes in Artificial Intelligence), Los Altos, California, 129–157.

Jacobs, Paul S. (1985). "A knowledge-based approach to language production." Technical Report No. UCB/CSD 86/254, Computer Science Division (EECS), University of California, Berkeley.

Jakobsen, Roman and Halle, Morris (1956). *Fundamentals of Language*. Mouton.

Johnson, Mark (1980). "A philosophical perspective on the problems of metaphor." In *Cognition and Figurative Language*, edited by Richard P. Honeck and Robert R. Hoffman, 47–68, Erlbaum Associates.

Katz, Jerrold J. (1964). "Analyticity and contradiction in natural language." In *The Structure of Language: Readings in the Philosophy of Language*, edited by J.A. Fodor and J.J. Katz, 519–543, Prentice-Hall.

Lakoff, George and Johnson, Mark (1980). *Metaphors We Live By*. Chicago University Press.

Lees, R.B. (1960). "The grammar of English nominalizations." *International Journal of American Linguistics*, 26(3).

Lenat, Douglas B. and Guha, R. V. (1990). *Building Large Knowledge-Based Systems: Representation and Inference in the CYC Project*. Addison-Wesley.

Levin, Samuel R. (1977). *The Semantics of Metaphor*. John Hopkins University Press.

Loewenberg, Ina (1975). "Identifying metaphors." *Foundations of Language*, 12:315–338.

Malgady, R. and Johnson, M. (1976). "Modifiers in metaphors: effects of constituent phrase similarity on the interpretation of figurative sentences." *Journal of Psycholinguistic Research*, 5:43–52.

Martin, James H. (1990). *A Computational Model of Metaphor Interpretation*. Academic Press.

Martin, James H. (1988). "A computational theory of metaphor." Technical Report No. UCB/CSD 88/465, Computer Science Division (EECS), University of California, Berkeley.

Martin, James H. (1985). "Knowledge acquisition through natural language dialogue." In *Proceedings, 2nd Annual Conference on Artificial Intelligence Applications*, Miami, Florida.

Matthews, R. J. (1971). "Concerning a 'linguistic theory' of metaphor." *Foundations of Language*, 7:413–425.

Miles, J. (1967). *Style and Proportion*. Little, Brown and Co.

Mish, Frederick C. (1986). *Webster's Ninth New Collegiate Dictionary*. Merriam-Webster Inc.

Modiano, Nicole (1986). "Two procedures for sense projection." Unpublished ms., Computing Research Laboratory, New Mexico State University.

Van Noppen, Jean-Pierre; De Knop, S.; and Jongen, R. (1985). *Metaphor: A Bibliography of post-1970 Publications*. Amsterdam Studies in the Theory and History of Linguistic Science (Series V: Library and Information Sources in Linguistics), Volume 17, John Benjamins Publishing Company.

Ortony, Andrew (1980). "Some psycholinguistic aspects of metaphor." In *Cognition and Figurative Language*, edited by Richard P. Honeck and Robert R. Hoffman, 69–83, Erlbaum Associates.

Ortony, Andrew (1979). "Beyond literal similarity." *Psychological Review*, 86:161–180.

Percy, Walker (1954). *The Message in the Bottle*. Farrar, Strauss, and Giroux.

Quillian, M. Ross (1968). "Semantic memory." In *Semantic Information Processing*, edited by Marvin Minsky, 216–270, MIT Press.

Perrine, Lawrence (1971). "Psychological forms of metaphor." *College English*, 33:125–138.

Reddy, Michael J. (1979). "The conduit metaphor — a case of frame conflict in our language about language." In *Metaphor and Thought*, edited by Andrew Ortony, 284–324, Cambridge University Press.

Reddy, Michael J. (1969). "A semantic approach to metaphor." In *Chicago Linguistic Society Collected Papers*, Chicago University Press, 240–251.

Richards, Ivor A. (1936). *The Philosophy of Rhetoric*. Oxford University Press.

Russell, Sylvia Weber (1976). "Computer understanding of metaphorically used verbs." *American Journal of Computational Linguistics*, Microfiche 44.

Schank, Roger C. (1975). "The structure of episodes in memory." In *Representation and Understanding*, edited by Daniel G. Bobrow and Allan Collins, 237–272. Academic Press.

Schank, Roger C. (1973). "Identification of conceptualizations underlying natural language." In *Computer Models of Thought and Language*, edited by Roger C. Schank and Kenneth M. Colby, 187–247. W.H. Freeman.

Shibles, Warren (1971). *Metaphor: An Annotated Bibliography and History*. The Language Press.

Stern, Gustaf (1968). (first published in Sweden 1931) *Meaning and Changes of Meaning*. Indiana University Press.

Tourangeau, Roger and Sternberg, Robert J. (1982). "Understanding and appreciating metaphors." *Cognition*, 11:203–244.

Tversky, Amos (1977). "Features of similarity." *Psychological Review*, 84:327–352.

Ullmann, Stephen (1962). *Semantics: An Introduction to the Science of Meaning*. Basil Blackwell & Mott Ltd.

Waldron, Ronald A. (1967). *Sense and Sense Development*. Andre Deutsch.

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

Weiner, E. Judith (1984). "A knowledge representation approach to understanding metaphors." *American Journal of Computational Linguistics*, 10:1–14.

Wheelwright, P. (1962). *Metaphor and Reality*. Indiana University Press.

Wilks, Yorick A. (1978). "Making preferences more active." *Artificial Intelligence*, 11:197–223.

Wilks, Yorick A. (1975b). "An intelligent analyzer and understander for English." *Communications of the ACM*, 18:264–274.
Wilks,Yorick A. (1975a). "A preferential pattern-seeking semantics for natural language inference." *Artificial Intelligence*, 6:53–74.
Wilks, Yorick A. (1973). "An artificial intelligence approach to machine translation." In *Computer Models of Thought and Language*, edited by Roger C. Schank and Kenneth M. Colby, 114–151, W.H. Freeman.
Wilks, Yorick A. and Fass, Dan C. (In press). "The preference semantics family." To appear in *Computers and Mathematics with Applications*, Special Issue on Semantic Networks in Artificial Intelligence.
Winston, Patrick H. (1980). "Learning by creatifying transfer frames." In *Artificial Intelligence: An MIT Perspective, Volume 1*, edited by Patrick H. Winston and Richard H. Brown, 347–376, MIT Press.
Yamanashi, Masa-aki (1987). "Metonymic interpretation and associative processes in natural language." In *Language and Artificial Intelligence (Proceedings of an International Symposium on Language and Artificial Intelligence held in Kyoto, Japan, 16–21 March 1986)*, edited by Makoto Nagao, 77–86, Elsevier Science Publishers.

Glossary of Main Terms

Anomalous relation: a semantic relation indicated by a violated preference and the absence of a relevant analogy [see Section 3].

Assertion: a word sense-based contextual constraint in which semantic information is imposed onto the local context of a word sense [Section 3].

Collation: a process that discriminates the semantic relation(s) between two word senses by matching the sense-frames for the word senses [Section 4].

Literal relation: a semantic relation indicated by a satisfied preference [Section 3].

Metaphor: a trope in which one entity is used to view another entity to which it bears a partial resemblance [Sections 2–7].

Metaphorical relation: a semantic relation indicated by a violated preference and the presence of a relevant analogy [Section 3].

Metonymy: a trope in which one entity is used to refer to another that is related to it [Sections 2–7].

Metonymic relation: a semantic relation indicated by failure to satisfy a preference and the presence of one or more conceptual relationships like PART-WHOLE [Section 3].

Preference: a word sense-based contextual constraint in which semantic information restricts the local context of a word sense [Section 3].

Screening: a process that resolves lexical ambiguity by choosing among semantic vectors on the basis of rank orderings among semantic relations and a measure of conceptual similarity [Section 4].

Semantic relation: the basis of literalness, metonymy, metaphor, etc.; found by evaluating lexical semantic constraints in sentences [Section 3].

Semantic vector: a data structure that represents semantic relations by recording the matches produced by collation [Section 4].

Sense-frame: a framelike data structure that represents a word sense and is composed of other word senses having their own sense-frames [Section 4].

Sense-network: a semantic network of word senses formed from information in sense-frames [Section 4].

Source: the lexical item in a semantic relation whose contextual constraint(s) are enforced [Section 3].

Target: the lexical item in a semantic relation on which contextual constraint(s) are applied [Section 3].

Trope: the technical term for a nonliteral figure of speech, e.g., metaphor, metonymy, simile, understatement (litotes), overstatement (hyperbole), and irony [Section 1].