An Integrated Model of Lexico-Semantic and Conceptual Processing for the Treatment of Natural Language Metaphors

Tony Veale,

School of Computer Applications,

Dublin City University,

Glasnevin, Dublin,

Ireland.

Email: Tonyv@compapp.dcu.ie

Mark Keane,

Dept. Computer Science

O'Reilly Institute,

Trinity College,

Dublin 2,

Ireland.

ABSTRACT

Research into the computational properties of the human metaphor faculty can be categorized according to two different perspectives: the lexico-semantic approach, most favoured by linguists, which sees metaphor as a predominantly linguistic phenomenon and thus attempts to model it as such; and the conceptual approach, which sees metaphor as a deep-seated cognitive agency that systematically organizes much, if not all, of our conceptual architecture. Both of these perspectives have brought different algorithmic and representational apparatus to bear on the task of metaphor interpretation and generation, but each is incomplete in itself. The lexico-semantic approach provides a good account of how different word combinations can give rise to creative figurality, while the conceptual approach explicates the deep structural mechanics of metaphor and analogy, connecting both into the breadth of world knowledge. This paper presents an integrated view of metaphor that unifies the lexico-semantic and conceptual approaches, explaining how metaphor is governed both by word and world knowledge. Emerging from this synthesis is a new role for metaphor in the workings of case grammars and the issue of case / filler compatibility.

Keywords: *Metaphor, Analogy, Lexical-Semantics, Conceptual Models, Natural Language Processing, Case Grammars*

An Integrated Model of Lexico-Semantic and Conceptual Processing for the Treatment of Natural Language Metaphors

Introduction

Any reference to "*natural language metaphors*" will most likely seem redundant to many linguists: since metaphor is primarily a linguistic phenomenon, surely our theories should treat it as such, and by definition have something to say about its grounding in natural language? In recent times, however, the most energetic work on metaphor has proceeded outside the sphere of purely linguistic interest, concentrating instead on the cognitive and conceptual basis for a putative metaphor faculty (typified by Lakoff and Johnson, 1980; Weiner 1984; Johnson, 1987; Hofstadter and Mitchell, 1988; Holyoak and Thagard, 1989; Way, 1991; Barnden 1992,1998; Indurkhya, 1992; Fauconnier and Turner, 1994; and Keane, Ledgeway and Duff, 1994). This movement signals a shift away from traditional linguistic concerns of sentence form and word meaning to the use of conceptual content and world knowledge.

This cognitive pursuit has been motivated by the uncontroversial view that the human metaphor faculty is primarily concerned with the generation and the creation of conceptual appositions, apt juxtapositions which direct the reader to view one domain through the conceptual lens of another. However, a theory of metaphor which concerns itself only with the interpretation of conceptual juxtaposition captures just one aspect of human metaphoric cognition, inasmuch as it ignores those higher-level structural aspects of utterance analysis that lead to the recognition of such juxtapositions in the first place. Since linguistic metaphor exploits creative wordplay to communicate novel ideas, the lexical and syntactic strokes of this wordplay deserve equal consideration to the composition of the ideas involved. But too often, metaphor analysis limits itself to the simplicity of the *"A is (like) B"* copula form (see for instance Weiner, 1984), conveniently side-stepping the vexing fact that most verb-centred metaphoric utterances evoke a considerable range of different juxtapositions. A metaphoric theory with aspirations to fullness must thus provide an algorithmic basis for how the structure of the metaphoric utterance—both syntactic and semantic—may be coaxed into giving up its underlying pattern of conceptual appositions, before these appositions can themselves be analyzed. In this paper we consider this interplay between the structure of the utterance and the structure of the underlying conceptual domains described by the utterance. This will lead us to consider, in broad terms, the relation between syntactic and semantic analysis, and the strategic nature of different metaphor comprehension processes.

Overall, this paper highlights the need to incorporate two different views of system knowledge into metaphor analysis: these views are i) *generalized lexico-semantic structures*, which distill experience into abstracted rule-like schemata and which are used to impose structure upon an utterance, and ii) *contingent memory-based knowledge,* which is used to elaborate and flesh-out such structures with real-world facts. The first of these resources is traditionally labeled the "dictionary", while the latter is commonly referred to as the "encyclopaedia" (e.g., see Eco, 1984). It is our contention in this paper that dictionary knowledge is best applied top-down, whilst encyclopedic knowledge should emerge, bottom-up, from semantic memory, with both meeting at a middle ground we term a *Scaffolding* structure. We thus demonstrate the benefits accruing from an integrated comprehension process that addresses the complementary concerns of how an utterance is to be structurally analyzed to yield a framework of conceptual associations, and how these associations are to be interpreted relative to the mass of background experiential knowledge possessed by the interpretative agent.

The present model can be seen both as a work of Cognitive Science (CS) and of Artificial Intelligence (AI). As a piece of AI, it is designed to demonstrate the computational benefits and economies of scale deriving from an integrated approach to metaphor comprehension, showing that integration is both computationally feasible and inferentially attractive. As a piece of CS, it is designed to reflect known psychological constraints on metaphor. For example, the model is designed to be consistent with the *total time constraint* on metaphor comprehension; namely, that given appropriate contextual cueing a metaphor should take no longer to process than the equivalent literal statement (see Hoffman and Kemper, 1987; Gerrig, 1989).

This paper observes the following logical structure: after presenting the linguistic and conceptual motivations for this work, as supported by a critique of the failings of some specific non-unified models, we proceed to consider the specific instantiation of an integrated architecture from which our conclusions have been drawn. The benefits of fusing two complementary perspectives on language and metaphor in this way are then considered, with a demonstration of how metaphor can bidirectionally mediate between the lexical and conceptual levels during the resolution of structural ambiguity. This mediating role is given computational expression by demonstrating the potential for using an existing model of structural metaphor as a case-filling mechanism in theories of case grammar.

Motivating Principles of an Integrated Approach to Metaphor

In many respects, lexical semantics is simply a convenient contrivance of the lexicographer, one that allows the morass of situational knowledge about a word's meaning to be distilled into a bounded and compartmentalized semantic packet (see Wilks, 1988 for a similar view). However, such discrete representations of word meaning are designed to contain enough information—generally of an analytic nature—to allow an aggregate semantic structure to be derived compositionally from a given string of words. Overall then, a lexical semantics allows us to decentralize the contents of long-term semantic memory, distributing and indexing it across the lexical spectrum of a language where it can be used most directly for text comprehension.

For these reasons, most approaches to the treatment of textual metaphors tend to be lexico-semantic in orientation, as lexical semantics provides a clear channel between word meaning and metaphor interpretation (see for instance Wilks, 1978; Aarts and Calbert, 1979; Fass, 1988; Weber, 1988; Russell, 1976, 1992). But equally clearly, lexical semantics does not offer a complete picture of utterance meaning. For while lexical semantics may yield a serviceable first-cut, it mostly fails to accommodate outside forces of interpretation that are placed on a metaphor by experiential and contingent knowledge of the real world. For instance, one might employ a metaphor to portray the precarious state of democratic Russia as "*A drunken bear whose body shows the scars of a quadruple heart-bypass*", or similarly, state that "*The Russian bear is but a heart-attack away from staggering back to Communism*". Such metaphors exploit for their full interpretation not only established relationships between Russia and bears, between heart-attacks and disasters, and between by-pass operations and regulatory laws (all of which one might expect to find in a lexico-semantic representation), but also contingent knowledge regarding the specific health worries that currently plague Boris Yeltsin. In this latter respect metaphor interpretation sometimes calls for creative metaphor generation if an allusion to the real-word is to be appreciated.

Conceptual models, on the other hand, do not concern themselves with the semantic relation between words and meanings, and are thus free to operate wholly within the domain of conceptual knowledge. This liberation from lexical concern is quite central to the success of conceptual models, inasmuch as the quantity of world knowledge that may potentially bear on the interpretation of a given metaphor would overwhelm any efforts to compartmentalize it into discrete word-indexed bundles. Nevertheless, as Wilks (1975) argues, a model which maintains tight relations between the lexical and conceptual levels is more likely to be semantically grounded than one which is not, and for this reason Wilks bases the primitives of his (1975) preference semantics model on words of the English language. Though apparently a circular argument, it appeals to intuition inasmuch as one can validate that his primitives have not been invented from whole cloth (a claim often leveled at the primitives of Schank, 1975), but derive a provenance from the language which they are designed to represent. That is, lexical semantics at least ensures a link between a machine's representation of a concept, and how people actually seem to use that concept.

Since words are used to *'refer'* in the real world, a clear mapping between the lexical and conceptual level serves to ground concepts in words, granting concepts *'referring power'* in the world also. Though this mapping between words and concepts is usually an ambiguous one, words do at least provide prima facie evidence for the mental reality of particular concepts. In contrast, invented concepts that defy explication in simple lexical terms (e.g., in words rather than paragraphs) tend to diminish the credibility of the models in which they are exploited. Since the lexical and conceptual levels of an integrated model must interact at a common ground, the structure of one will serve to constrain the other. For instance, difficulties in representing an idea lexico-semantically will discourage the use of overly ad-hoc categories that do not reflect the way these concepts are communicated via language.

Immediate examples of such ad-hoc categorization can be found in Way (1991), who uses the concept THINGS-WHICH-BEHAVE-IN-A-MYSTERIOUS-OR-SINISTER-MANNER as a common ground for the metaphor *"Nixon is a Submarine"*, and in Gentner, Falkenhainer and Skorstad (1989), who use LIFE-AS-A-FLUID-PROCESS-THAT-FLOWS-FROM-THE-FUTURE-INTO-THE-PAST as a grounding for the Virginia Wolff metaphor *"She allowed life to waste away like a tap left running"*. These analyses exemplify the kind of strong rationalism, or mentalese, against which empiricists have always reacted, namely, the generation of ungrounded mental constructs wherever it is convenient to do so. Such constructs can abound when one conveniently confines oneself to the level of conceptual analysis, and thus does not have to motivate them at the level of lexical semantics, that is, at the level at which people should most visibly use such constructs. We argue that one should therefore theorize about both levels *at once*, in the context of a unified model.

As demonstrated by the work of Lakoff and Johnson (1980), lexical issues constrain conceptual issues to such an extent that one can exploit systematic variations of word usage as a reliable empirical basis for inferring general patterns of the human conceptual structure. Additionally, Rosch *et al.* (1975; 1976) note that those concepts which occupy the most natural level of human categorization—termed the 'basic level'—exhibit the most direct lexical mappings, being typically expressible in a single word. Since Rosch's evidence demonstrates that humans clearly consider basic-level concepts to be most useful in terms of naturalness and discriminatory power, it follows that semanticists should do so also, and exploit basic-level concepts wherever possible. A unified theory that balances the strains of lexical semantics with the stresses of conceptual representation will most likely meet then at a middle ground that is governed by Rosch's basic-level of categorization.

Psychological Motivations

Thus far we have stressed the computational aspects of our model and said little about the psychological motivation behind the work. Though the integrated model is in many respects more exploratory than psychological, it has been designed to be sensitive to important psychological constraints on metaphor processing, thereby explaining why constraints which seem psychologically real arise for computationally valid reasons. In particular then, the model is consistent with various theoretical proposals and empirical findings on metaphor.

Theoretically, the model is designed to meet the *total time constraint*; this constraint proposes that, given appropriate contextual cueing, a metaphor should take no longer to process than an equivalent literal statement. This constraint is supported by repeated findings from the psycholinguistics literature (see Hoffman and Kemper, 1987; Gerrig, 1989). We model this constraint by ensuring that our integrated model performs equivalent operations on each utterance, without necessitating the sentence to be pre-judged as either literal or figurative (see Lytinen *et al.*, 1992, for proposals on how monolithic or non-integrated models might handle this constraint).

Continuing in this theoretical vein, the model is consistent with several proposals on metaphor. For instance, the model's exploitation of domain systematicities means it is consistent with other contemporary views on metaphor such as those of Gibbs (1992) and Lakoff and Johnson, 1980). Following these views, our approach asserts that metaphors exist in complex systematic relationships to each other at a conceptual level, and that much of linguistic metaphor is a manifestation of deeper metaphoric schemas (such as the BODY-AS-CONTAINER; ANGER-AS-PRESSURE; MOOD-AS-ORIENTATION). The current approach acknowledges the deep conceptual roots of metaphoric language, and explicitly models the process whereby a novel metaphors can be interpreted as an elaboration of an established conventional metaphor. Our unified approach is also consistent with the category inclusion model of Glucksberg and Keysar (1990) and Glucksberg, McGlone and Manfredi (1996). In fact, the model coheres with both views despite evidence that these theories are not themselves mutually supportive (see, for instance, McGlone, 1996).

The category inclusion model claims that metaphors are a means by which existing categories can be broadened: to say *'T is a V'* is to say that T belongs to category or class of which V is typical. For example, the metaphor *'Jobs are jails'* places the concept JOB into the class of restrictive and claustrophobic situations, of which the concept JAIL is an exemplar. The process of category inclusion does *not* attempt to find an underlying metaphoric schema (such as SITUATION AS CONTAINER) against which to interpret the statement. The main conceptual mechanism described in this paper, Sapper, is not explicitly a model of category structure, and thus does not claim to directly implement the category inclusion theory. Nonetheless, we believe the model is consistent with Glucksberg's theory. Sapper attempts to establish rich, consistent mappings between the structure of two concepts (such as JOB and JAIL). In doing so, it generates a domain theory as to why one concept can be viewed as another (e.g., a Sapper theory of why a JOB can be JAIL would recursively point to the similarity between BOSS and WARDEN, RESPONSIBILITY and SHACKLE, and OFFICE and CELL). As such, Sapper too can be seen as a model of category extension, and thus, as a cognitive mechanism of category inclusion.

As well as positing a representational role for conventional spatial metaphors, the unified approach also employs another representational structure that is claimed to be cognitively real — the *conceptual bridge*. Bridges are the means by which the Sapper model has a representational effect on long-term memory, thus allowing the model to make claims about how metaphor can create new associations amongst concepts, and not merely exploit existing associations (see Camac and Glucksberg (1984) for empirical arguments as to why this is a characteristic aspect of human metaphor use). The notion of a bridge is a compelling one, and plays a key role in other models of conceptual creativity, from Hofstadter *et al.* (1995) to Hummel and Holyoak (1996) (where the latter employ the notion in the guise of *mapping connections*).

Linguistic Theories of Metaphor

Generally speaking then, we can distinguish linguistic theories of metaphor from knowledge-based theories by the relative concentration of each on the importance on linguistic form versus conceptual content. Linguistic theories tend to focus on the syntactic expression of a metaphor, while viewing interpretation as largely a problem of lexical semantics. Knowledge-based theories tend to focus on the conceptual models underlying a metaphoric opposition, and often take as given this opposition without the need to perform a structural analysis of the utterance. Thus, the linguistic models of Aarts and Calbert (1979) and Russell (1976/1992) locate their computational models within those processes that extract a semantic structure from a given surface structure, while knowledge-base models such as that of Weiner (1984), Gentner, Falkenhainer and Skorstad (1989), and Holyoak and Thagard (1989) either assume the surface structure of the metaphor to have a simple templatized form or assume that the apposition has a priori been extracted from the text. In this section we present a brief critique of the first two linguistic models, while also pointing to certain deficiencies in Weiner's knowledge-based model, in support of our argument for an integrated linguistic/conceptual approach.

The Aarts and Calbert approach

Building upon a semantic marker theory of compositional semantics originally proposed by Katz and Fodor (1964), Aarts and Calbert (1979) provide an extension of the basic marker scheme that is intended to model the metaphoric use of adjective-noun combinations within the linguistic framework of generative semantics.

The Aarts and Calbert extension to this scheme compartmentalizes the marker-based (or *markerese*) definition of each word into four different feature sets. The *primary feature* set holds the core semantic markers called for under the original Katz and Fodor scheme. The *secondary feature* set is posited to hold markers which represent *frozen* metaphoric extensions of the word—essentially dead metaphors. Thus, while the word "red" is primarily defined as a color (with selectional constraints for +PHYSICAL and +VISIBLE), its secondary compartment also contains a marker representation for "communist" (with the selectional constraint +HUMAN). The *generative feature* set then contains markers that allow for the word to be used in a different, but related sense, as is the case in synecdochal and metonymic usages. However, the most interesting compartment, from a metaphor processing point of view, is the *transfer feature* set, which contains semantic markers that may be copied across into the representation of the tenor word[1] in metaphoric situations. These transfer features are intended to capture the abstract, *skeletal* meaning structures that underlie everyday word use. Thus, for instance, the words "Big" and "Heavy" are each tagged with the transfer feature *+INTENSE*, enabling their model to assign the same interpretation to the conventional metaphors "Big Problem" and "Heavy problem", whereby a transfer process causes the tenor word "Problem" to be tagged as *+INTENSE*.

This is relatively superficial approach to metaphor, one that fails to account for the systematicity that is so apparent in even the most mundane of metaphors (e.g., see Lakoff and Johnson 1980). Consider the figurative adjective:noun combination “narrow minded”, in which Aarts and Calbert argue there occurs a transfer of a negative connotation from the transfer set of “narrow” to the word

---

1 We employ here the terminology of Richards (1936), who refers to the target of a metaphoric description as the *tenor*, and the source of that description as the *vehicle*. Thus, in a metaphor of the

“mind”. Now, while it is of course true to say that *narrow-mindedness* is indeed a negative and undesirable quality, being indicative of a poor mind, the true workings of this metaphor work around not some abstract notion of a negative scale, but more plausibly, around a conventional *focus* metaphor that is culturally deep-seated. At a knowledge-level, rather than lexico-semantic level, narrowly focused channels of communication can be seen to produce a limited intellectual range. One might also say that narrow-minded people consider the range of allowable alternatives using a narrow search-light, and that narrow minds are therefore less illuminating.

Likewise, other scale adjectives, such as “small” and “shallow” are also applicable here, but with subtle nuances of meaning that are entirely beyond the scope of a flat marker representation. Small-minded is suggestive of intellectual pettiness, while shallow-minded implies superficiality of thought. Other modifiers in the same metaphor family, springing from the same conceptual model of mind, include “open”, “closed”, “dense”, “thick” and “empty”. As Barnden (1992) discusses in detail, there are a variety of inferentially deep metaphors of mind that one might bring to bear on these descriptions. A computer model, like a person, should be allowed some latitude for subjectivity in their application, but the point is that a systematic and non-superficial analysis is required. This is not to say that use of skeletal representations such as a positive/negative connotation scale are wrong-headed, rather that their use should be complemented by some deeper, knowledge-based resource.

Pathological interpretations can abound in a system where deep systematicities are short-circuited by shallow representations. Consider for instance “Blind justice”; from the perspective of Aarts and Calbert, “blindness” is an example of dysfunction and thus “blind” should transfer a negative connotation, say *-FUNCTION*, to “decision”. However, the intended interpretation of “blind justice” is positive, not negative; there is a sense in which “blind” can indeed mean “without foresight or insight”, but here it equally well mean “impartial”. The solution to this ambiguity is not to be found in lexical semantics, but in a knowledge-based understanding of the impartiality of the law. Likewise, a markerese interpretation of “Heavy heart” will paint over the systematic metaphor that relates grief to a burden that must be carried. Indeed, such an interpretation is likely to associate a positive marker such

---

form *A is (like) B*, A may be referred to as either the tenor or the target of the metaphor, while B is referred to as the vehicle or source.

as +*INTENSE* with “heavy” (which is suggestive of *moreness)*, and thus pathologically view a heavy heart as a desirable possession.

The Russell approach

As a maturation of her earlier structure-mapping approach to metaphor (see Russell 1976), a clever elaboration of the Aarts and Calbert approach is advocated by Sylvia Weber Russell, in which concepts are described in terms of partitioned sets of semantic markers (see Russell 1992). In her new account, concepts are defined in terms of both specific, domain-dependent *non-extensible* features, and abstract, domain-independent *extensible* features. In the terminology employed by Aarts and Calbert (1979), these partitions correspond to the *primary* feature set and the *transfer* feature set respectively, inasmuch as the primary or non-extensible set defines the core specifics of a concept, while the transfer or extensible set defines the abstract essence of the concept, a transferable/extensible *distillation* that allows the concept to be projected into other domains.

Having created an initial conceptual representation of a metaphoric utterance, using case information associated with verb definitions and other conceptual structures, Russell's system, known as MAP, then proceeds to project this structure out of the ill-suited literal domain and into a more apropos target domain. The basis for this mapping is an abstract structure created from the extensible features of the concepts involved. Thus, while non-extensible features are *variable*, since they are expected to change when moving to a new domain, extensible features should be considered *invariant*, as it is they which provide the common ground between the tenor and vehicle domains.

A structure/case matching algorithm is used to move between representations in different domains. For example, in interpreting the metaphor "*He gave her his opinion*", a representational movement occurs between the first cut representation of the utterance (what Martin (1990) calls the *'primal representation'*)—a physical transference reading—to a mental communication reading, the former being semantically anomalous but the latter being semantically proper. The mechanics of this domain transfer have been clarified substantially in Russell (1992)—her earlier approach being open to claims that it hard-coded particularly generative conventional metaphors into the match algorithm, thereby forfeiting generality (and cognitive reality) for limited effectiveness (see Martin, 1990 for one such critique). Her latest approach employs a markerese scheme to represent the extensible or domain-independent content of different actions (for instance, she links the CD primitives PTRANS, ATRANS and MTRANS of Schank (1975) via a common extensible feature +TRANSFER), and thus the connection between *mappable* actions in different domains is made at the representational, rather than the algorithmic, level. Ultimately, however, Russell's approach suffers from a similar lack of conceptual depth as the marker model of Aarts and Calbert, exhibiting a dependency upon lexical semantics (i.e., the dictionary) that diminishes the importance of world knowledge (i.e., the encyclopaedia).

This distinction of dictionary / encyclopaedia lies at the heart of the difference between linguistic and knowledge-based approaches to metaphor. The limitations of dictionary (lexico-semantic) versus encyclopaedic approaches are well illustrated with an example from Russell. Consider for instance the MAP interpretation of the metaphor "*The news torpedoed his hopes*", which MAP comprehends in terms of an abstracted structure "*The news CAUSEd his hopes NOT TO BE*" that is generated using lexico-semantic rather than world knowledge. To be sure, this is a very limited interpretation of TORPEDO, an object which rarely causes anything to cease to be (for instance, the *Bismarck* and the *Lusitania* still exist), but whose main purpose is to blow large holes in the hull of marine vessels, thus causing them to sink. In the vast majority of cases, vessels which are sunk in this way cease to be functional, and thus a torpedoed hope is, prototypically speaking, one that demonstrates no possibility of fulfillment. However, even this inference chain from *Torpedo—cause→Loss-of-Buoyancy—effect→ Obsolescence* is not itself canonical, as buoyancy is itself a conventional English metaphor used to denote a mental state of happiness and general optimism about the world. It is this conceptual underpinning that makes a torpedo such an apt metaphoric device for the destruction of one's hopes. An equally believable account is that the prototypical usage of the verb "To Torpedo" casts its object in the role of a doomed naval vessel, with the effect that the negative effects of the action (submersion, obsolescence) become applicable, via analogical mapping, to the object in question. Just as a torpedo blows a hole in a ship, thus diminishing its structural integrity, bad news can diminish the integrity of a theory (which is, in essence what a hope is — an optimistic theory of the future) by removing some of its supporting facts. In fact, one can argue that "to torpedo" is simply an extension of the conventional metaphor "to poke a hole in an argument". Regardless of the metaphor's provenance, however, this chain of analogical reasoning, based as it is in world knowledge, is forever outside the scope of a purely lexical account of metaphor.

Knowledge-Based Models

Admittedly, the exemplar we cite to characterize the linguistic approach, that of Aarts and Calbert (1979), is an extreme one. There are many computational models of metaphor that model the interpretation process from surface syntax to deep meaning without demonstrating the obvious failings of this case study. For instance, the *meta5* system of Fass (1988), a model which is perhaps easier to pigeonhole as a lexico-semantic rather than purely conceptual approach to metaphor, provides an interesting mechanism (and a variant of preference semantics known as *Collative Semantics*) for tackling metonymy and case ambiguity as they occur in metaphoric language. The related *Metallel* system of Iverson and Helmreich (1992) employs an additional spreading activation process over a knowledge-base to resolve more complex metonymies. In a different vein, the *MIDAS* system of Martin (1990) embodies the systematicity notion of Lakoff and Johnson, employing a case memory of past metaphors against which it attempts to understand new metaphors as extensions of old, while the *Dynamic Type Hierarchy* (DTH) model of Way (1991) employs a rich ontology of world concepts to interpret metaphors in terms of common superclasses.

The Weiner approach.

Weiner (1984) provides a knowledge-based treatment of metaphor, one which posits the notions of *salience imbalance* (see Ortony 1979) and *prototypicality* (see Rosch *et al.* 1975; 1976) as being central to the metaphor interpretation process. Using a semantic network representation (the KL-ONE model of Brachman, 1979) to model domain knowledge, Weiner augments the relations of this representation with numeric salience measures. These measures indicate to the metaphor system which properties of the vehicle should most plausibly be transferred to the tenor, following Ortony's theory of salience imbalance (in which highly-salient properties of the vehicle should be transferred if these properties possess a lower salience in the tenor). Essentially, Weiner's approach represents a pragmatic synthesis of existing ideas on metaphor, borrowing *sub-theories* where necessary from Ortony and Rosch. Weiner's model may be considered then an *implementation guide* to certain theories of metaphor, inasmuch as she provides ample detail on how a knowledge-base manager such as KL-ONE can be used to support predicate transfer and salience imbalance.

However, while providing a well-motivated theory of feature transfer and emphasis in metaphor, Weiner's theory is essentially feature-based rather than structure-based, for the most part avoiding complex issues of structural mapping in metaphor. For instance, it was argued in our critique of Russell that metaphor interpretation often necessitates the reconciliation of two structurally rich domains (e.g., hopes and naval vessels) by finding appropriate points of similarity around which an analog mapping can be constructed (e.g., structural integrity). This mapping is essentially a domain theory, one that relates not only the tenor and vehicle concepts, but other concepts from both domains that hold the same relative positions. For example, the metaphor *"Surgeons are butchers"* may well rely on some feature-based understanding of surgeons and butchers, but it almost certainly relies on a related understanding of scalpels, cleavers, operating theatres, and animal slaughter as well. What is required then from a knowledge-based theory is an appreciation of the systematic interplay of structures from both domains that occurs during metaphor interpretation. Such a theory should, unlike Weiner's, be capable of analogically transporting an arbitrarily complex conceptual structure from a domain in which an inference is unsupported into one in which it is.

Since Weiner takes the practical step of considering only metaphors of the form *"A is (like) B"*, her emphasis on the representation of world knowledge is clearly achieved at the expense of linguistic representation. But while identity metaphors are a useful and common class of metaphors, much of the content of a metaphor is carried in its linguistic form. To appreciate the full semantic import of a metaphor then, a unified approach that gives equal emphasis to linguistic and conceptual aspects of meaning must be adopted. For instance, before the key appositions of a metaphor can be comprehended, they must initially be extracted from the surface form of the linguistic utterance. In this respect, linguistic form acts as a top-down, shape-inducing pressure that guides the bottom-up knowledge-based processes of comprehension.

The Iverson and Helmreich approach: Metallel.

Perhaps the model that comes closest to meeting the goals of unified processing as discussed in this paper is the Metallel system of Iverson and Helmreich (1992). Metallell, which is an extension and reworking of the Meta5 system of Fass (1988), combines two weak inferencing methods – path-finding in semantic networks (in the style of Charniak, 1983), and lexical template-matching (in the style of Wilks, 1978) – to construct an interpretation of a sentence in which the processes of word-sense disambiguation, metonymy resolution and metaphoric interpretation are performed in a coherent, knowledge-based framework.

Metallel employs a process of *path-finding* in a network model of lexico-conceptual knowledge to determine chains of semantic relations that link verbs to their stated fillers. All things being equal, Metallel prefers shorter paths to longer, more tenuous pathways, and will thus prefer word senses of ambiguous words that lead to the most direct relationship between the verbs and nouns of a sentence. For example, the concepts CAR and DRINK can be connected via the chain of concepts CAR—*isa*→VEHICLE—*isa*→MACHINE—*isa*→DEVICE—*isa*→PHYSICALOBJECT←*isa*—ORGANISM←*isa*—ANIMAL←*agent*—DRINK, which actually represents two pathways that converge at a common node, PHYSICALOBJECT. Likewise, the filler concept GASOLINE can be linked to verb concept DRINK via the converging pathways GASOLINE—*isa*→LIQUID←*isa*—POTABLEDRINK←*patient*—DRINK. So given a sentence such as *'My car drinks gasoline'*, the fillers CAR and GASOLINE can be reconciled with the semantic preferences of DRINK, namely ANIMAL and POTABLEDRINK, by using this pathfinding process to converge on a set of common superordinates. In this way, Metallel uses marker passing or spreading activation to model a view of metaphor interpretation that has been representative of approaches from Aristotle (1982) to Wilks (1978).

Because metonymy is a phenomenon whereby semantic associates are used to stand as lexical *proxies* for a given concept (e.g., using 'Washington' to stand in place of the U.S. government), Metallel's path-finding approach naturally accommodates metonymic uses of concepts without recourse to specialized rules of metonymic construal, as used in Fass (1988). Typically, the path-ways that reconcile a verb with a metonymic filler (as in a 'ham sandwich' requesting his bill) will be longer than those needed to reconcile strictly literal fillers, but shorter than those needed to reconcile metaphoric fillers. Based on path length (effectively conceptual distance) alone, Metallel should tend therefore to prefer metonymic interpretations than metaphoric ones for non-literal sentences.

However, Metallel also uses default world knowledge, or micro-facts, stored in the lexical definition of its word senses, to impose coherence on a potential interpretation. If a given interpretation, as reached via an analysis of semantic pathways, coheres with a known fact about the domain in question, that interpretation is preferable to Metallel than an alternative interpretation arising out of shorter pathways that does not reflect world knowledge in this way. For instance, Metallel knows that a CAR consumes GASOLINE, and that PEOPLE consume COFFEE, but not vice versa. Thus, a metaphoric interpretation of *'My car drinks gasoline'* will cohere with world knowledge, while a metonymic one will not. Likewise, a metonymic interpretation of *'That car drinks coffee'*, where it is the passengers that are seen as consumers, is preferable to a metaphoric one that would receive no support from the established wisdom in Metallel's knowledge-base.

Moreso than its predecessors, Metallel reflects a necessary synthesis between word and world knowledge. But it is not a complete synthesis. In its movement away from pre-established rules of construal to handle metonymy, such as ARTIST FOR ART-WORK (e.g., 'Bill plays Mozart'), Metallel works from first principles and thus also eschews the use of pre-established metaphors. Lakoff and Johnson (1980) demonstrate that much of metaphor use is actually a reworking or elaboration of highly conventionalized metaphors, which we treat as conceptual building blocks in their own right. The idea of a car drinking gasoline is an example of well-worn personification metaphor in which machines are viewed as people, with goals, needs and intentions of their own. Metallel's path-finding mechanism has no explicit mechanism for recruiting such established metaphors to reconcile the metaphoric uses of words with their intended meanings. Likewise, Metallel is not designed to process system metaphors in which many mutually consistent mappings are created. For instance, one might say that *'My car drinks decaff'* to convey the notion that one's car uses unleaded gasoline. This interpretation requires a mapping between LEAD and CAFFEINE, which in turn supports a mapping between UNLEADED and DECAFFEINATED. The more complex example *'My wife's VW drinks decaff, but my Porsche drinks espresso'* elaborates this system metaphor further, to map ESPRESSO to a HIGHPERFORMANCEFUEL, and thus, my PORSCHE to a HIGHPERFORMANCECAR. Mappings like these require a coherence mechanism different from that employed by Metallel, one that can synthesize different but consistent pathways arising from the same statement. In effect, this type of coherence requires a model of analogical mapping of the kind conventionally used to model the conceptual, rather than the linguistic, uses of metaphor.

What More Needs To Be Done?

Ultimately, then, we argue that each of these knowledge-based accounts do not go far enough. We claim that what is indicative of a truly knowledge-based approach is the power to plumb the depths of memory in response to a theory proposed by a higher-level of analysis, to construct conceptual hypotheses of sufficient richness and complexity, while nevertheless staying within the bounds of computational tractability. These higher-levels, such as syntax, then have the freedom to compare the richness of the conceptual structures supporting alternate theories of attachment and scope, to choose the interpretation that is most consonant overall with the contents of long-term memory. The structures we shall be considering throughout this paper have an analogical richness that other models we have mentioned, with the possible exception of Martin's MIDAS system, do not.

Architecture of an Integrated Approach to Metaphor

We propose an open, tiered architecture for integrating the mechanisms necessary for processing figurative language. The architecture is tiered in the sense that it posits two distinct layers of abstraction at which the meaning of an utterance may be considered. Broadly speaking, this dichotomy reflects that which exists between the dictionary (word knowledge) and the encyclopaedia (world knowledge) as complementary resources in understanding language. More specifically, the primary tier of this architecture concerns the lexico-semantic processing of the utterance, for it is at this level that a working representation of utterance meaning is compositionally constructed from the dictionary semantics of its constituent words. (Researchers interested in sound symbolism and morphological blending, as in "Bollywood" = "Bombay" + "Hollywood", may need to posit an even earlier, phonological tier at which word meanings are constructed). The second tier contains a range of elaboration processes that extend and flesh-out the initial structure created by the first, thereby generating a complete interpretation of the utterance.

This tier is *'open'* in two different senses. Firstly, it is open in a resource sense, since we posit that a wide variety of meaning resources—not simply the lexico-semantic definitions of a dictionary—will be brought to bear here. Some elaboration processes may indeed be lexical in that they represent expectations about certain word collocations and usages, but others will be primarily conceptual and pragmatic in nature, bringing the encyclopedic range of world knowledge to bear on an interpretation. These processes are also *open* in an architectural sense, since they may vary from one speaker community to another, or from one system implementation to another, and thus may not be exhaustively enumerated in any single model.

The specific instantiation of this general architecture that we employ in the current work is the *Conceptual Scaffolding* model of Veale and Keane (1992a,b). This model of meaning has as its starting point the work of Lakoff and Johnson (1980), who argue that no metaphor can be comprehended, or even adequately represented, without regard to its experiential basis in the physical world. They claim therefore that this experiential basis has to form part of a metaphor system's basic knowledge representation. Martin's MIDAS system (1990) partially addresses this requirement, by incorporating explicit knowledge of core metaphors into the system’s design, but the issue of experiential grounding remains for the most part implicit in this *metaphoric knowledge* approach. In contrast, the first tier of a Conceptual Scaffolding harnesses a range of spatial or 'localist' structures, which are explicitly grounded in physical reality, and which are extensible into increasingly more abstract domains of discourse. We use *localism* in its linguistic sense here (see e.g., Anderson, 1971; Lyons, 1977), to refer to the empiricist school of language that views the structure of even the most abstract ideas to be ultimately founded in spatial intuition (this is the sense in which we use the term throughout; we do not use it to denote the class of neural network representations where neurons correspond directly to concepts). For instance, the class of emotions can be broadly organized along localist lines: e.g., we frequently speak of 'being *high'*, 'cheering *up'*, 'coming *down'* and *'falling into* a bad mood'.

The second tier uses a variety of elaboration strategies to build around and (mostly, to replace) the underlying scaffolding of the metaphor that is built during this first, lexico-semantic stage of processing, thereby reworking the skeletal shell of the metaphor’s thematic structure with richer, domain-specific conceptual relations. Of the processes working within this tier we single out here two wholly conceptual mechanisms, which we dub Sapper and Scout. The Sapper model of Veale *et al.* (1993,1994,1995,1996a,b,1997) is a structure-mapper in the tradition of the SME model of

Falkenhainer *et al.* (1989) and the ACME engine of Holyoak and Thagard (1989), and works at the level of conceptual representations to determine the appropriate points of connection between different knowledge structures. Scout is a complementary mechanism of structure-retrieval, one that is responsible for locating in memory the appropriate representations for Sapper to work with. Scout is necessarily employed then in utterances where the tenor of a metaphor is not explicitly stated, but merely alluded to, as is frequently the case in poetry.

In effect, the elaboration stage builds around and replaces the scaffolding structure to yield as full an interpretation of the utterance as possible, given the current state of domain knowledge possessed by the system. This process is not necessarily complete, or immediate, for a lack of domain knowledge can result in an incomplete interpretation. However, any holes in the interpretation are effectively *plugged* by the underlying scaffolding, which maintains the shape of the intended meaning structure while further construction is awaited. As such, this stage is open ended, supporting an ever increasing number of different elaboration strategies, and providing an ideal framework for the acquisition of new conceptual knowledge. The scaffolding representation can maintain the intended association of ideas, and so bridge a gap in our conceptual repertoire (thereby forming a useful basis of inference), until the appropriate conceptual structure is eventually acquired.

Consider the conventional metaphor "*Mary caught the flu*", which describes the spread of an intangible entity (influenza symptoms) in terms of the transfer of a physical object or projectile. For the majority of comprehenders, the spatial basis of this metaphor, as expressed in the lexical-semantics of "To Catch", will be elaborated via a localist schema such as CONCEPT-ASSOCIATION-AS-COLLOCATION. That is, the comprehender's conceptual model of MARY will most likely become associated with the conceptually salient features of INFLUENZA. This allows the negative connotations (or *'implicative complex'* in the terms of Black, 1962) of INFLUENZA to become associated, through feature transfer, with MARY. In contrast, comprehenders who possess additional domain knowledge about the physical nature of contagion and viral infection may not employ a metaphorically-grounded schema such as this, but might instead understand the utterance physically in terms of a model of germ transfer.

We argue that it is frequently possible to perform naïve or *folk* inferencing directly at the level of localist structure, since conceptual collocations often underlie physical collocations inasmuch as more than mere geographical information is conveyed. For example, the import of *"Only Nixon could go to*

China" is more political than geographical, and one appreciates this import via a juxtaposition (or collocation) of the conceptual structures NIXON and CHINA, and indirectly, via that of RIGHT-WING and LEFT-WING. Metaphors can thus be interpreted both at a basic localist level which is meaningful inasmuch as it supports naive inference, and at an elaborated level that connects into the relevant domains of discourse. The degree of elaboration corresponds to the one's competence in these domains, and the richness of domain knowledge that a comprehender can bring to bear. Thus, for the utterance "Mary caught the flu", a doctor is less likely to reason at the folk-level via localist schemas such as CONCEPT-ASSOCIATION-AS-COLLOCATION, and more likely to use specific domain knowledge about the physical transfer of germs.

In the remainder of this section we examine the key elements of both levels of the architecture, beginning with a consideration of scaffolding structures and their construction, before discussing two of the most prominent conceptual mechanisms brought to bear in the elaboration of these structures, Sapper and Scout. We will then consider how these mechanisms are actually used in different elaboration strategies to interpret a variety of linguistic forms.

The Lexico-Semantic Level: Scaffolding Structures

The tiered architecture divides the meaning construction process into two phases, lexical and conceptual: the first uses lexical-semantics to construct a *scaffolding* structure which captures the essential themes of the utterance, in a vague and somewhat sweeping form, while the second constructs an elaborated meaning using pragmatic and encyclopaedia knowledge. In this respect scaffolding structures follow very much in the vein of Langacker's *Cognitive Grammar* approach (1991), which was itself initially termed *Space Grammar*. This scaffolding structure interrelates the elements of the utterance through the use of spatial metaphors as basic meaning constructors. These spatial metaphors are also used to define the basic lexical semantics of words, allowing a scaffolding structure to be compositionally constructed from the localist semantics of the words in an utterance. For instance, the word "Happy" can be defined as an upward orientation of the concept MOOD, while "Marriage" can be defined as a spatial collocation of people (a "coming together"). These localist definitions exhibit many parallels with the *Lexical-Conceptual Semantics* of Jackendoff (1983; 1990), though as we will see, a defining property of scaffolding structures is their potential for future domain enrichment or elaboration.

The localist elements of a scaffolding structure can be elaborated to yield spatio-physical interpretations (i.e., to represent actual spatial relations) or figurative interpretations (e.g., using a schema such as ABSTRACT STATE AS LOCATION). Indeed, a scaffolding might be elaborated in a number of complementary ways, and still arrive at the same broad meaning, the sentence "Mary caught the flu" being a case in point. Here the meaning can be interpreted as the non-physical coming together of ideas and their associated features (as in "The film caught the mood of the era perfectly"), or as the spatial coming together of physical entities (i.e., the human body and germs). The important point to note here is that no such interpretation occurs when the scaffolding is constructed. The scaffolding structure is *indeterminate*, yielding neither a literal or metaphoric interpretation until it is later elaborated. It thus captures the common ground upon which literal and metaphoric uses of the same concepts are based. Thus, a scaffolding cares not whether one can physically *catch* a disease as one would a ball, since this is an issue for future elaboration processes.

In this respect, the scaffolding idea is analogous to a recent movement toward underspecification in formal semantics, as exemplified by the *Quasi-Logical Forms (QLFs)* of Alshawi (1990) and the *Underspecified Discourse Representation Structures (UDRSs)* of Reyle (1995). As with UDRS (but unlike QLFs) a scaffolding structure supports a certain amount of direct inferencing, based on experiential spatial intuitions of the world (e.g., disconnection from an upwardly-oriented state results in a downward orientation, while connection with an upwardly-oriented state results in an upward orientation), which may later be elaborated in non-spatial terms. And as with QLFs, scaffoldings can often be used as a semantic transfer level in machine translation, without recourse to further specification, due to the near-universal nature of the spatial constructors employed (see Veale and Collins, 1996).

Meaning constructors of the conceptual scaffolding.

The scaffolding model currently defines a small but productive set of constructors for defining the lexical semantics of words. These operators may be supplemented by additional non-spatial terms, but in general, we claim that the spatially-oriented core set should act as the backbone of most lexical definitions. **Connect** and **Disconnect** embody the highly general CONCEPT-ASSOCIATION AS

COLLOCATION schema (as in “I *sunk to* his level” and “He *found* a solution”) and a variety of social RANK-AND-DISTANCE schemas (as in “*Father* of the A-bomb” and “A *sister* company”). The **Contain** and **Release** constructors are a container-based elaboration of this basic idea (since containment implies collocation), and can be seen at work in sentences like “how do I *get into* EMACS?” and "He *fell into* a depression". **Compare, Contrast** and **Identity** are further elaborations of the basic collocation schema, and underlie metaphors of the form "A is (like) B” (e.g., "The rainforests *are* the lungs of the Earth"). The orientation constructors **Up** and **Down** support the highly productive GRADABLE-OPPOSITION AS VERTICAL-ORIENTATION schema, at work in “His mood is *high* today”. Finally, the non-spatial constructors **Attempted-Cause**, **Actual-Cause** and **Enabling-Cause** represent the possible causal relations between other scaffolding constructors. A graphical notation for these constructors is illustrated via examples in Figure 1.

Figure 1 around Here, “Basic Scaffolding Examples”

The basic ideas captured in a scaffolding structure are intended to reflect our localist intuitions about high-level concepts. For instance, ACME's receivership is represented as a downward financial movement to reflect its falling bank balances, share-prices and social status (i.e., DOWN IS BAD). IBM's divorce from Microsoft involves a spatial disconnection because this is the way we frequently conceptualize divorce (e.g., we talk of *'marital separation'* and *'getting back together again'*). Likewise, IBM's marriage to Apple involves spatial connection because we typically conceptualize marriage as a coming together (e.g., 'who would *sunder* what God *hath brought together*?'). 'Spatial' in this context means 'spatial metaphor', of course; the scaffolding structure is neutral with respect to its final interpretation, whether that be physical (e.g., IBM moves as far away from Microsoft's Seattle campus as possible) or abstract (e.g., IBM severs business links with Microsoft). These specific interpretations are the responsibility of the elaboration phase of comprehension.

In many respects the Up and Down constructors of the Scaffolding are merely spatial analogs of the +/- polarity markings exploited by Aarts and Calbert (1979) through to Weber (1988). However, the key to the Scaffolding’s expressive power lies in the ability to combine these constructors causally to build still richer structures. For instance, Figure 1 illustrates how these constructors typically combine to form relatively complex meaning structures, using the graphic notation developed in Veale and Keane (1992). Our cultural traditions of representing family organizations as trees demonstrates that spatial notions such as these underlie much of our thinking about kinship terms (e.g., '*up from* Fred we find his uncle *on* his mother's *side*). It is not surprising then—given the systematicity that is so often present in metaphor—that spatial metaphors are exploited to structure other social concepts as well. So not only can these constructors represent the underlying meaning of strictly physio-spatial actions, they are equal to the representation of higher-order abstractions as well, such as corporate rivalry (a disconnection), corporate mergers (a connection), divorce (an emotional disconnection) and marriage (an emotional connection).

Note that these spatial operators are intended to augment, rather than usurp, the role of traditional lexico-semantic meaning operators. Though much of the lexically-expressible meaning of an utterance can be captured by a spatially-oriented semantics, the Scaffolding model is nevertheless open to other, non-physiospatial constructions if such constructions exhibit a broad applicability to the meaning of a large class of words in an analytic fashion. For example, the temperature connotations of colour terms (e.g., WARM versus COLD), and the gender connotations of certain country names (e.g., FATHERLAND versus MOTHERLAND), exhibit sufficient applicability. We believe that lexico-semantic constructions that do not fall into this category are best represented at the conceptual level, where they can be tempered by world knowledge.

As outlined in Veale & Keane (1992), scaffolding representations also provide a solid basis for systematic inferencing from a metaphor. If one claims "*IBM has left behind the slump of the '80s*", it is possible to restructure this disconnection from a figurative downward position as a connection with an upward position, which is subsequently elaborated as a state of greater financial prominence. Likewise, the raw spatial nature of the scaffolding throws any unsystematicities into sharp focus; for instance, the utterance "*After Windows 3.0, Microsoft fell into an all-time high*" is incoherent, positing a simultaneous downward and upward position for Microsoft.

To provide a more knowledge-based perspective, Figure 2 presents a scaffolding analysis of *"Chomsky rebuilt modern linguistics from the ground up"*, combining a frame representation with this spatial notation to illustrate the utility of scaffolding as a meaning substrate.

Figure 2 around Here, "Scaffolding/Chomsky Example"

This example illustrates the need for combined top-down and bottom-up interpretation processes: the scaffolding representation imposes a broadly thematic conceptual organization on the utterance *from above* (i.e., using general schematic knowledge), capturing the underlying causal structure while indicating which concept juxtapositions should be considered further *from below* (i.e., using contingent world knowledge). Modern linguistics is subsequently viewed as an edifice/construction, possessing the theoretical equivalent of the concepts *structural integrity* (logical soundness) and a *ground* (logical foundations). These mappings are constructed bottom-up *from memory*, as detailed in the description of Sapper in a later section.

Construction details: how a scaffolding is built.

Construction of the conceptual scaffolding for a given utterance is a compositional process, one that combines the partially instantiated scaffolding structures in each word's lexical definition to create an overall scaffolding graph. These definitions are similar to the *attribute-value-matrix*, or AVM, definitions used in linguistic models of unification such as PATR (see Shieber, 1986) and HPSG (see Pollard and Sag, 1987), or lexico-conceptual models such that of Jackendoff (1983; 1990). For example, the lexical definition of DIVORCE will contain an uninstantiated Disconnect operator that links two unbound variables, which in turn co-reference with the Agent and Patient roles of the verb. So when used in a sentence such as *'Microsoft divorced IBM'*, these variables become bound via the compositional process of structure building (i.e., as the roles of the verb are filled with the given noun-phrases). The scaffolding approach thus requires no special algorithms that are not already a matter of convention in parsing models such as Shieber's PATR (*Parsing and TRanslation*).

Consider as an example the utterance *'IBM fell into a coma'*. The lexical structures associated with each of the words in this utterance are illustrated in Figure 3 (the structures shown are necessarily simplified for purposes of presentation). Note how the scaffolding structure for the word 'IBM' is simply the token IBM, and not a complex meaning structure. The consensus knowledge one normally associates with IBM—that it is a huge multinational computer company, inventor of the PC personal computer, and home to assiduous dark-suited salesmen who one counted Ross Perot among their number—is stored as world knowledge in the semantic network representation of long-term memory. Scaffoldings provide structured indices into this knowledge, but rarely constitute this knowledge itself.

Figure 3 around Here, "IBM fell into a coma."

Variables are illustrated in Figure 3 as boxed numbers, or *co-reference indices*, to use the standard notation of unification grammar. As words are combined into phrases by the syntactic analyzer, local scaffolding structures are unified to provide more complex (and more fully instantiated) scaffoldings. Grey lines in Figure 3 indicate the flow of the compositionality process. *'A'* is combined with *'Coma'* to yield the scaffolding structure *down(Health)*, while 'IBM' yields the simple scaffolding IBM directly; when these structures become respectively bound to the INTO and PATIENT roles of the verb *'To Fall'*, they become, by virtue of co-referencing, bound into the scaffolding of the verb to produce the fully instantiated structure *contain(down(*HEALTH*),* IBM*)* – in other words, IBM has entered a state that corresponds to a downward turn of health.

Note that the local scaffolding in *'To Fall'* which disconnects the PATIENT from its original SOURCE state (as in 'IBM fell *from grace'*) does not become fully instantiated in the current utterance, and thus does not contribute to the overall scaffolding structure. Note also that the Patient role of *'To Fall'* contains a partially specified scaffolding structure, *down(* 1 *)*, which means that only fillers that possess a downward orientation can fill this role. In this way the unification mechanism enforces a basic form of metaphoric coherence. Infelicitously mixed metaphors, such as *'IBM fell into a recovery'*, will thus not be parsable (it is debatable, of course, if a system should enforce coherence at this level; that is a matter for future research).

It is important to remember that because scaffolding structures are *pre-interpretative*, compositionality of construction does not imply compositionality of interpretation (as it does, in contrast, in the approach of Lytinen *et al.* 1982). As we shall argue in a later section, many interesting metaphors require that the overall scaffolding structure be interpreted holistically, rather than via an incremental interpretation of its parts.

Conceptual Components

We now consider two related conceptual mechanisms, Sapper and Scout, which work not at the level of linguistic expression, but at the level of conceptual representation. Sapper provides a model of how concepts from different domains can be analogically or metaphorically reconciled by determining coherent graph isomorphisms between their representational forms. As such, Sapper also provides a mechanism of category extension, since such reconciliation shows how a domain representation can be extended to include the representation of a foreign concept. Scout is a complementary model of structure retrieval, which determines for a given domain concept a counterpart from another domain that will yield a rich Sapper isomorphism. Working together, Sapper and Scout serve the conceptual needs of a variety of elaboration strategies.

Sapper.

The Sapper framework, as described in Veale and Keane (1996, 1997), is a hybrid symbolic and connectionist model, embodying a basic philosophy which views the interpretation of novel metaphor as a process of connectionist *bridge-building*. From the Sapper perspective, metaphor comprehension involves the construction of new linkages, or bridges, between different domains of discourse; these bridges then serve to bind the *analogous* correspondences that underlie a metaphor. The novelty of a metaphor may be measured then by the extent to which it adds to the structure of its semantic network, as the metaphor is accommodated into the system's long-term memory. A metaphor which gives rise to a host of new bridge links, causing the semantic terrain of long-term memory to be reshaped, should be considered more novel and creative than one that adds none, or one that merely reinforces bridges that are already in place from previous metaphors.

Memory organization in Sapper.

Sapper views long-term memory as a symbolic graph, or semantic network, in which nodes represent concepts, and where arcs represent semantic relations between these concept nodes. Activation energy can spread across these relational arcs from concept to concept, causing each node in turn to be stimulated and made contextually active (see Quillian, 1968; Collins and Loftus, 1975; Charniak, 1983).

New cross-domain bridges are laid down along previously established, but *dormant* and non-activation carrying, bridges between concepts. These dormant bridges have been created a priori by special rules of structural similarity that automatically structure the contents of memory. Dormant bridges therefore determine the possible routes along which metaphoric creativity can arise. Though by definition creativity should involve a novel, almost anarchic combination of concepts, it is still a constrained process which stretches rather than breaks convention, for general structural constraints are what divide creativity from nonsense. (For instance, even the music of Mozart was ultimately constrained by the structure of the chromatic scale.) Sapper embodies the idea that these basic constraints on creativity can be modeled in terms of simple rules that enforce a local consistency in the structure of semantic memory. Sapper's rules of structural similarity thus analyse the organization of memory for local consistencies of structure, and subsequently augment memory with new (but dormant) connections on the basis of these consistencies.

Dormant connections represent merely *plausible*, rather than fully *established*, semantic relations, and are thus not operative carriers of activation. Sapper uses the controlled propagation of activation (i.e., Quillian's notion of spreading activation) from the vehicle concept node to mark those concepts in the vehicle domain that might possibly contribute to the interpretation of the metaphor. The goal of metaphoric mapping in Sapper is to find those dormant connections linking the tenor domain to the vehicle domain that should eventually be used as sites on which to build new conceptual bridges. The act of building a new, activation-carrying bridge can have serious ramifications for the structure of memory. As new bridges are built between domains, the number of potential activation pathways between those domains increases, and thus, so do the interference effects due to spreading activation. From the perspective of inference processes that rely on spreading activation as a guide (e.g., determining deductive closure), such domains will have effectively moved *closer together*, and become more interactive (in the sense of Black, 1962).

Memory usage during metaphor interpretation.

Figure 4 depicts such a scenario during the interpretation of the metaphor *"Composers are Generals."* When waves of competing activation meet at opposing ends of a bridge (either dormant or active) then that bridge forms the basis of a potential mapping between the corresponding elements of the tenor and vehicle domains. At this point, Sapper determines whether the same chain of conceptual relations (such as PART, CONTROL, ATTRibute, etc.) can be used to link the tenor side of the bridge to the originating tenor node, and the vehicle side of the bridge to the originating vehicle node. If so, Sapper has identified a converging pair of isomorphic semantic pathways that can be placed in analogical alignment to generate a partial interpretation of the metaphor. For example, the chain of relations CONTROL→PART→CONTROL can be used to link the concepts COMPOSER to MUSICAL-INSTRUMENT via MUSICIAN and ORCHESTRA, while the same chain links GENERAL to MUSKET via SOLDIER and ARMY. Thus the bridge linking MUSICAL-INSTRUMENT to MUSKET acts as a grounding, in terms of literal similarity, for a possible partial interpretation that maps not only these two concepts together, but that also maps ORCHESTRA to ARMY and SOLDIER to MUSICIAN.

Since dormant linkages are laid down on the basis of local consistencies of structure, the mapping hypothesis stage is thus, in a strong sense, driven by both literal similarity and higher-order structural constraints. And because activation originates at the concept nodes of the tenor and vehicle, the connectionist phase effectively combines the memory-retrieval process (sometimes called *'base filtering'*) with that of the generation of mapping hypotheses. A fuller view of semantic memory during the interpretation of this metaphor is illustrated in Figures 6 and 7 following.

Figure 4 around Here, "Activation Flow"

Continuing with this example, Sapper considers a range of partial interpretations in response to the juxtaposition of COMPOSER and GENERAL. For instance, by mapping the pathway COMPOSER—*control*→ORCHESTRA—*part*→PERCUSSION—*part*→DRUM to the pathway GENERAL—*control*→ARMY—*part*→ARTILLERY—*part*→CANNON, an even richer partial interpretation can be created around the bridge linking DRUM to CANNON. Sapper uses a seeding process to rank these partial interpretations according to their richness (i.e., the number of distinct cross-domain mappings that each contains), allowing the richest to be combined to form an overall mapping that is both coherent (i.e., in terms of a one to one mapping of tenor and vehicle concepts) and systematic (i.e., each mapping is justified by higher-level mappings). Eventually then, Sapper determines the mappings of Figure 5 as being its favoured interpretation of the COMPOSER AS GENERAL metaphor. Dormant bridges that correspond to any of these mappings are subsequently awakened by Sapper, to become full carriers of activation in future metaphor interpretations. Memory thus records the processing of the metaphor in a way that can be exploited by future metaphors.

Figure 5 around Here, "Composer:General Mappings"

In the next section we focus the rules of structural similarity that lay down dormant bridges in long-term memory in the first place. For the moment though, we assume such bridges as given and present an overview of the main Sapper interpretation algorithm in Figure 6. The algorithm follows a greedy seeding procedure which is similar to that used in Forbus and Oblinger (1990), which, while not guaranteeing optimal results, provides an acceptable polynomial-time bound for its operation.

Figure 6 around Here, "Sapper Algorithm"

Conceptual similarity.

The number specified in square brackets to the left of each mapping in Figure 5 is a measure of the perceived similarity, in the range -1 and +1, of the related concepts *after* the metaphor has been comprehended. This measure, described more fully in Veale (1995), combines a metric for both the literality similarity of the related items (e.g., drums and cannons are loud, heavy and both go *Boom!* quite often), and the higher-order similarity that is now seen to exist between the two (e.g., the relation between DRUM and CANNON supports the second-order mapping PERCUSSION AS ARTILLERY, which in turn supports the mapping ORCHESTRA AS ARMY).

Sapper employs two distinct constructor rules to augment long-term memory with dormant conceptual bridges — the *Triangulation Rule* and the *Squaring Rule*. In essence, these rules compile into memory the top-down knowledge that will later be necessary to infer new systematic cross-domain bindings. Because this knowledge is automatically pre-compiled into the memory network, the activation-based process of memory retrieval and hypothesis formation is thus spared the necessity of performing a global structural analysis of its hypotheses, as the very existence of a dormant bridges implies a local systematicity of conceptual structure. The general form of these rules is illustrated in Figure 7.

Figure 7 around Here, "Triangulation and Squaring"

The *Triangulation* rule is invoked whenever two concept nodes share a common association or superclass, as is the case with the concepts MUSICAL-INSTRUMENT:MUSKET and HAND-HELD, DRUM:CANNON and LOUD, BATON:SABRE and LONG, BATTLE-THEATRE:CONCERT-THEATRE and THEATRE, and STAVE:CONTOUR and LINE. New dormant bridges are thus established between the concepts MUSICAL-INSTRUMENT and MUSKET, DRUM and CANNON, BATON and SABRE, BATTLE-THEATRE and CONCERT-THEATRE, and STAVE and CONTOUR. In essence, the triangulation rule is a formalization of the principle underlying the plan recognition model of Hendler (1989), in which two high-level concepts can be seen as plan analogs if they share one or more task-specific micro-features. For instance, an antique letter-opener can be recognized as a workable substitute for a knife in a KILLING-PLAN, being an object that is sharp enough to accomplish the task at hand (MURDER), yet one which—unlike a knife—will not arouse suspicions going through airport customs. The same intuition is used in the LISA model of Hummel and Holyoak (1996), whereby two high-level concepts can be viewed as analogous if they relate to the same set of active low-level semantic features.

The *Squaring* rule is a *second-order constructor* that acts upon the bridges established by the triangulation rule (whether dormant or newly active) to lay down further dormant bridges between even higher-level concepts. For example, tracks laid down by the squaring rule underlie the bridges ORCHESTRA:ARMY, PERCUSSION: ARTILLERY, MUSICAL-INSTRUMENT: MUSKET, and MUSICAL-SCORE:BATTLE-PLAN. In effect, Sapper employs the squaring rule to ensure that any low-level similarities that are discovered by the triangulation rule are percolated up to higher-level concepts in a structurally coherent fashion.

Triangulation and Squaring can be expressed more formally as follows:

- **Triangulation:** *If memory already contains two linkages L_{ij} and L_{kj} of semantic type L forming two sides of a triangle between the concept nodes C_k, C_i and C_j, then complete the triangle and augment memory with a new dormant bridge B_{ik}.*

- **Squaring:** *If B_{jk} is a conceptual bridge, and if there already exists the linkages L_{ij} and L_{lk} of the semantic type L, forming three sides of a square between the concept nodes C_i, C_j, C_k and C_l, then complete the square and augment memory with a new dormant bridge B_{il}.*

The Squaring rule is a means by which bridges can be built upon bridges, each new bridge extending the inter-domain *reach* of the last; in this way Sapper accounts for the phenomenon of *domain incongruence*—the non-literal sharing of attributes across domains (see Tourangeau and Sternberg, 1981). For instance, the incongruence between LOUD-CLOTHING and LOUD-NOISE can be explained via a squaring relation with the concepts GARISH and NOISY, which are in turn connected by triangulation via the shared properties INTENSE and UNPLEASANT. From Sapper's perspective, creative metaphor generation is thus a matter of recursively percolating local, almost banal similarities up to a level at which they seem to be global incongruities.

Figure 7 around Here, "Composers and Generals"

Scout: A model of analog retrieval.

On the flip side of the metaphor coin lies the generation of apt, high quality metaphors. This process complements that of metaphor interpretation, involving the search for a suitable vehicle structure from semantic memory to satisfy the system's communication goals regarding a particular tenor. To this end, several analog retrieval models have been proposed in the literature, such as the MAC/FAC model (*Many Are Called but Few Are Chosen*) of Law, Forbus and Gentner (1994), which is the complementary retrieval component of the SME *(Structure-Mapping Engine)* of Falkenhainer, Forbus and Gentner (1989), and the ARCS model *(Analog Retrieval by Constraint Satisfaction)* of Holyoak, Gochfeld and Nelson (1990), which similarly complements the ACME model *(Analog Constraint Mapping Engine)* of Holyoak and Thagard (1989).

While primarily of conceptual rather than linguistic concern, the issue of analog retrieval does impinge on the interpretation of metaphoric utterances. This is because the meaning of a metaphor is often pragmatically carried by what is implied rather than what is explicitly stated, meaning a metaphor system must look past the meaning of individual words to consider the implications of the scenario pictured collectively by those words. For instance, a metaphor that portrays modern Russia as a drunken bear is all the more powerful if one recognizes it as a dual satire on Boris Yeltsin. Likewise, in comparing Chomsky to an architect one compares Chomskyan linguistics to a piece of modern architecture: to fully appreciate the metaphor, one may thus feel inclined to consider what kind of architecture best typifies Chomskyan linguistics. The reader may believe that the modernist designs of

Frank Lloyd Wright, with their almost mathematical lines and minimalist elegance best typify the central ideas of generative linguistics.

Or consider another example, *"Hackers regulate the circulatory system of the computer world"*. A system can recognize this as a metaphor because computer worlds do not possess circulatory systems, but just what are we referring to when we use "circulatory system" in this sense, and how exactly do we perceive the role of a hacker? In underspecified metaphors such as this, it clearly becomes an issue of interpretative significance to be able to probe the appropriate areas of long-term memory for hypotheses as to the potential references of metaphoric allusions, and to perform this memory search in a tractable fashion. In the hacker example, this necessitates the use of spreading activation from the concept matriarch HACKER to find an analogous structure that relates to, and explains the intended allusion of, the concept CIRCULATORY-SYSTEM.

We turn describe such a model of spreading-activation, named Scout, which exploits Sapper's use of conceptual bridge links between domains. Scout is essentially a local-memory search algorithm, traversing all semantic contours within a horizon H of the probe concept in search of potential cross-domain bridges. Once found, each bridge is pursued into another domain, whereupon Scout's original path is replayed in reverse, to reach (if possible) a potential analog concept for the probe. Thus, like Sapper, each path leading from the tenor to a bridge-head must be mirrored by an isomorphic pathway leading from the opposing bridge-head to the vehicle. Spreading activation, which is a weak unstructured heuristic, is given structure by the constraint of isomorphism, causing misleading bridges which do not connect isomorphic pathways to be ultimately rejected. As in Sapper, each such pairing of pathways serves as a partial interpretation for a potential metaphor between the given tenor and the newly located vehicle. Scout collects all partial mappings that are constructable with the horizon H, and chooses that interpretation which is the most mapping rich (i.e., composed from the longest paths). The intuition here is that the vehicle domain which yields the richest partial mapping should also yield a rich overall mapping. Scout thus chooses an analog concept based on the likelihood that it will generate a rich metaphoric interpretation. Scout's loop-based search process is illustrated in Figure 9.

Figure 9 around Here, "Hackers are Surgeons"

As illustrated in Figure 9 , the concept schema for SURGEON provides a very good fit for that of HACKER, while accounting for the allusion to "circulatory system" in the original expression of the metaphor by introducing the sub-metaphor INTERNET: CIRCULATORY-SYSTEM. Each sub-metaphor generated in this fashion may be necessary to the interpretation of further allusions in future statements in the same narrative; for instance, the ostensibly literal statement *"The circulatory-system is slowly becoming clogged with pollutants"* can be interpreted—after a suitable analog has been found for POLLUTION—as referring to the pollution of the internet with pornography. In short, analog retrieval is a necessary process in the interpretation of allusion, and is thus necessary to an integrated model of metaphor.

Figure 10 around Here, "Scout Algorithm"

The Scout algorithm is presented in more formal terms in Figure 10. Note how this algorithm is essentially a one-sided version of the Sapper algorithm of Figure 6, inasmuch as it attempts to construct a symmetric loop of isomorphic semantic pathways (crossing a single metaphoric bridge) linking the tenor to a vehicle concept. Unlike Sapper, however, Scout does not know what the vehicle domain is in advance. This makes Scout less constrained than Sapper, since whenever it follows a bridge into another domain, its path explorations there are not guided by a prior process of spreading activation (from a known vehicle) that delimits what concepts are potentially relevant.

Clearly, the more bridges leading out of the given tenor domain, the greater the search space that Scout must explore. But this is cognitively desirable, since bridges represent mapping opportunities with other domains, and thus, domains rich in bridges should have greater metaphoric potential that domains with few bridges. We should also expect then that generating metaphors from such productive concepts will be a more cognitively labor-intensive process that generating metaphors from concepts which provide little scope for comparison with others.

Pragmatic goals provide the most decisive constraints on Scout's search space. For instance, one can mark certain concepts in memory as forming mandatory points in a loop; thus, one can mark CIRCULATORY-SYSTEM as being a necessary part of the analog retrieved for the Hacker probe. This guarantees that the final metaphor will yield a mapping for the concept CIRCULATORY-SYSTEM, since the partial interpretation produced by Scout corresponds to the seed mapping that will eventually be used by the Sapper algorithm. Likewise, one can mark (via spreading activation) an entire space of concepts from which the vehicle analog must be drawn. For instance, if searching for an analog for MOTHER-IN-LAW, one might prime all concepts within H relations of the concept PSYCHOLOGICAL. This would help ensure that only analog structures that primarily describe the mental (rather than physical) aspects of mother-in-laws are retrieved, while simultaneously reducing the size of Scout's search-space.

Elaboration Strategies

We envisage two broad classes of elaboration strategy that can apply to a scaffolding structure. The first of these we term *projective strategies*, as these concern the transfer of semantic features from one concept to another, in much the same way as posited by the lexico-semantic approaches of Aarts and Calbert, and Russell, described earlier. The second class of strategy, which we term *mapping strategies*, employs Sapper and/or Scout to delve deeper into the structures of the concepts concerned, and attempts to establish an isomorphic alignment of these structures. Projective strategies are applicable when the concepts concerned share one or more functional or causal dimensions that allow salient features from one to be directly predicated to another, while mapping strategies must be used when the concepts concerned are too far removed in imaginistic terms to share features in this way. But once a structure-mapping is established, inferences can be drawn from the set of correspondences that is produced, and projective strategies can in turn be applied to these correspondences (which are conceptual collocations in their own right) if they are sufficiently close to share featural dimensions.

Projective strategies.

Many conventional collocations lend themselves to interpretation by projective strategies. For instance, *"Bill gave the walls some much-needed yellow paint"* establishes a collocation between WALL and YELLOWPAINT. Since A functional dimension of WALL is APPEARANCE (as well as STRENGTH and LOCATION), and a functional dimension of PAINT is APPEARANCE (in this case, YELLOW), the feature YELLOW can be transferred from YELLOWPAINT to WALL (as opposed to, say, the features FLUID and VISCOUS). Likewise, in *"Holyfield gave Tyson a punch"* the feature PAIN is a causal property of PUNCH that is transferred to TYSON as a result of the collocation established by the verb "To Give".

Projection can still be a valid strategy even when the collocated concepts do not share obvious descriptive dimensions, if a conventional metaphor can be found to mediate between the two. For instance, in *"IBM went into a coma"*, the concept COMA is lexically defined in terms of a downward orientation on the dimension HEALTH (see Figure 3), yet the concept COMPANY does not possess a HEALTH dimension. However, the conventional metaphor FINANCIALSTATUS AS HEALTH (i.e., *financial health*) allows this downward orientation to apply directly to IBM's financial status. For a detailed description of various projective strategies, see Veale and Keane (1992).

Mapping strategies.

In contrast, mapping strategies employ Sapper and Scout to establish a recursive set of correspondences between the elements that are collocated. In the simplest mapping strategy, a collocation *X*→O←*Y* prompts Sapper to align the conceptual structures of *X* and *Y*. However, people frequently employ collocations with subtle nuances that must be reflected by the strategies that interpret them. In general then, a different mapping strategy is needed for each such nuance, as described below:

Collocation Mapping Strategies for X→O←*Y*

Nuance: Two physical entities in the same physical location (spatial collocation) *Example*: Bill is at MIT now.

Strategy: This corresponds to the physical realization of collocation. The strategy needs to find some aspect of the conceptual structure of *Y* (MIT) with which *X* (Bill) can correspond. For instance, MIT contains both STUDENTs and LECTURERs, and BILL can map to either (if the system knows particular features of BILL, such as YOUTH, this may cause the mapping with STUDENT to be richer than that with LECTURER).

Spreading activation is spread from *Y* (MIT) to mark out the relevant elements of this domain. Then Scout is used to find an analog for *X* (Bill) in this domain (e.g., STUDENT or LECTURER). Sapper can then produce the full mapping.

Nuance: Two entities are identical (identity as referential collocation)

Example: Clark Kent is Superman; Ted Kaczynski is the unabomber.

Strategy: A composite conceptual structure is created to represent both *X* and *Y* together. Coreferenced entities will tend to have many literal similarities (i.e., Clark and Superman are both white, adult males located in Metropolis; Kaczynski and the unabomber are both luddites) with little analogical overlap. Thus, analogical interpretations that produce rich mappings (e.g., 'Clinton is Hugh Heffner') should typically be preferred to coreference interpretations.

Nuance: Two ideas share some governing idea or domain (domain collocation) *Example*: Psychologically, Mary is a Sherman tank!

Strategy: Spreading activation is used to delimit the domain of relevant concepts reachable from the topic (PSYCHOLOGY). Sapper is then used to generate a structure mapping between *X* (MARY) and *Y* (TANK), and those mappings not primed by spreading activation are rejected.

If the system knows little or nothing about MARY, an empty mapping will be generated, yet the metaphor is still meaningful. A general backup strategy in this case is to spread activation from *Y* (TANK) also, using a different marker color to distinguish it from the first process. Bridges (dormant or otherwise) that connect concepts marked by alternate colors thus represent domain incongruities (or sub-metaphors) that link PSYCHOLOGY to TANK, e.g., SARCASM AS GUNFIRE, THICKSKIN AS ARMORPLATING, OBSTINATE AS UNSTOPPABLE. These concepts can then be placed in association with *X* (MARY).

Nuance: Two entities are in analogical correspondence (analogical collocation) *Example*: Elephants were the tanks of ancient warfare.

Strategy: Sapper is used to generate a mapping between the conceptual structures of *X* and *Y*. If a topic domain is specified (e.g., ANCIENTWARFARE), this can be used as the basis of a spreading activation filter. Mappings which do not involve concepts marked by this activation are then rejected.

In the containment schema $X \rightarrow\Box\ Y$, the *landmark* concept Y acts as a container (metaphoric or otherwise), while the *trajector* concept X moves into this landmark (we use here the spatial grammar terminology of Langacker, 1991). As noted earlier, the idea of containment is a specialization of collocation, and this specialization introduces additional nuances of its own:

Containment Mapping Strategies for $X \rightarrow \square Y$

Nuance: An entity is spatially or temporally bounded by another (spatio-temporal containment)

Example: Bill is in GeorgiaTech now; In the Summer Bill works for McDonalds.

Strategy: The same strategy is used as for spatial collocation above. However, use of containment additionally implies that *Y* is a container (abstract or otherwise), with associated boundaries. If *Y* cannot be conceived as having boundaries (e.g., starting/ending points) then use of the containment metaphor should be considered infelicitous.

Nuance: An entity is in a certain abstract state (ideational containment)

Example: Bill is *in* a foul mood; Mary is *into* existentialism; Ted is *in* Sales now;

Bill *has* abandonment issues; Dick *has* Ebola (possession is containment)

Ted is *in* purgatory; Bob was *in* hell *at* UCSD.

Strategy: Most conventional state containments can be processed using projective strategies, if the states concerned have highly salient features that describe appropriate dimensions of the tenor (e.g., EBOLA affects the tenor's HEALTH dimension in a variety of graphic ways). Containments using metaphoric states such as HELL or PURGATORY can also be proceeded using mapping strategies, if sufficient conceptual structure is known about the tenor to support a mapping. For instance, if we know that TED lectured at HARVARD, then Sapper may be able to construct a mapping between TED and SINNER, PURGATORY and HARVARD, and possibly, YALE and HEAVEN (depending on Ted's belief structures).

Nuance: An idea is bounded by another (domains are containers)

Example: Kinematics *belongs to* physics; Hannibal*'s* tanks ...

Strategy: Spreading activation is used to mark out the contents of the *Y* container domain (e.g., HANNIBAL). Scout is then used to find an analog for the concept *X* (e.g., TANK) in this domain, by restricting its explorations to concepts previously marked by the activation process. In this case Scout finds ELEPHANT to be good Hannibalic counterpart to TANK. Likewise, the colliding billiard balls of KINEMATICS should align with the colliding sub-atomic particles of the PHYSICS domain.

Of course, the domain containment may be a literal (i.e., conventional) one, such as *"Hannibal's elephants"* or *"Rommel's tanks"*. In these cases, the spreading activation process from Y will have already reached and marked X. The concept X will thus be seen to be its own best analog in the domain of Y, and Sapper need not be invoked.

Complementary mapping strategies also exist for the negative counterparts of collocation (disconnection) and containment (release). As described in Veale and Keane (1992), these negative constructors can be interpreted by first interpreting their positive counterparts, and inverting the results. For instance, in *"Psychologically, Mary is* no *Sherman tank"*, the concepts SARCASM, OBSTINATE and THICKSKINNED are determined as per the domain collocation strategy, and then *disassociated* from Mary.

Since a system cannot know in advance which strategy will be successful, all strategies must be applied in parallel. The issue then becomes: if more than one strategy returns a sensible interpretation (e.g., a non-empty structure-mapping, with potential for feature projection), which strategy should be chosen? Though this question may be deeply influenced by pragmatic concerns (what are the speaker's goals, etc.? see Veale and Keane, 1994), we propose at base a simple decision metric. Simply, the strategy that causes Sapper (and/or Scout) to produce the richest structure-mapping, in terms of the number of correspondences generated, is deemed to be the strategy of choice.

A Worked Example

Given our discussion of the Scaffolding and Sapper/Scout models, the representational demands and algorithmic mechanics of both are thus quite different, yet each comprises a co-operative part of a combined whole. We consider here a worked example that illustrates precisely how these models interact throughout the course of metaphor comprehension. A metaphor which readily demonstrates the need for bipartite model of comprehension can be found in the play (and film of same) “Amadeus” by Peter Shaffer (see Shaffer, 1985). In act II, the villain of the piece, the embittered and jealous composer Antonio Salieri, opens the scene as the overture to Don Giovanni is played on stage. This is Mozart's darkest opera, and concerns the fate of a libertine (Don Giovanni himself) who is damned by his accuser, the old General, and cast into hell for his transgressions. Salieri immediately recognizes the parallels between the relation of Don Giovanni to his accusing father-figure, and the relation of the young Mozart to his now dead father Leopold (whom Salieri describes as *"A father more accusing than any in opera").* The opera is seen then as both a tribute and a desperate plea for forgiveness to a father who still controls his son's life from beyond the grave. One can sum up the mood of this scene with the metaphoric utterance *"Mozart constructed Don Giovanni for his dead father Leopold",* using a CONSTRUCTION AS COMPOSITION metaphor to suggest that the opera is intended as a shrine for the composer's deceased father.

The initial step in comprehending a linguistic metaphor is structural analysis, from which a deep-structural and semantic representation is obtained. In the pages to follow, a computer trace of this analysis is reprinted in italics, while a parallel commentary is interleaved in roman face.

Employing the language-processing functionality of the Twig / Zardoz systems (see Cunningham and Veale, 1991; Veale and Cunningham 1992; Veale and Keane, 1998), the system proceeds to create a case-frame representation, labeled CONSTRUCT-0, for the utterance. In this frame, the filler WOLFGANG-AMADEUS-MOZART (which the system assumes is the intended reference of the proper noun "Mozart") occupies the role of *Constructor,* while the filler DON-GIOVANNI occupies that of *Construction,* and LEOPOLD fills that of intended *Recipient*. However, the system also knows that the prototypical fillers (or semantic *preferences*, in the mold of Wilks, 1975) of the *Construction* and *Constructor* roles are EDIFACE and ARCHITECT respectively. The collected cases and fillers of a verb thus constitute a *prototypical definition* of the verb in the sense of Rosch *et al.* (1975,1976). A particular instantiation of a verb frame might agree completely with its prototype, and therefore be considered a most orthodox (or 'literal') use of the verb, while those instantiations which disagree in some or all respects will represent unusual or novel usages.

Logically speaking, while nothing precludes an entity from being both an architect *and* a composer, a closed world assumption does lead the system to a pragmatic *impasse*. Furthermore, a frame-based organization of memory reveals that the relationship PROFESSION has a filler valency of one; that is, an entity frame can store a single filler in the slot PROFESSION, and in the case of MOZART, this slot is known to be occupied by a different value (i.e., COMPOSER). The system thus notes a tension between the expected and given fillers of the frame, spurring the evaluation of two new metaphor schemas, COMPOSER AS ARCHITECT and OPERA AS EDIFICE, to reconcile the fillers with the expectations.

This tension is the only sop to the literal meaning position embodied in this approach to metaphor. The utterance represented by the case-frame CONSTRUCT-0 is not deemed in the least part anomalous or nonsensical due to this tension, rather it is seen as all the more interesting because of its perceived departure from prototypical norms. This prompts the system to *explain away* the tension by means of some theory in which composers are conceptually reconciled with architects, rather than *spirit away* the case-frame itself (e.g., to become COMPOSE-0) as is done in many *anomaly-driven* theories of metaphor (for instance, Wilks (1978) and Lytinen *et al.* (1992) advocate such an approach).

We now consider a system trace of the current metaphor:

(parse Mozart constructed Don Giovanni for his dead father Leopold)

```
Create semantic preposition LINK-CONSTRUCT-0-FOR-FATHER
```

```
 0   8  [2, 9] - V  :  MOZART   CONSTRUCT-0   DON-GIOVANNI
        LINK-CONSTRUCT-0-FOR-FATHER HIS-0 DEAD FATHER  LEOPOLD

Surface Structure:

 (S     (NP MOZART)
        (V P   (VP   (V CONSTRUCT-0) (N DON-GIOVANNI))
               (PP (P LINK-CONSTRUCT-0-FOR-FATHER)
                     (NP (PRO HIS-0)
                           (NP  (NP   (ADJP (ADJ DEAD))
                                           (N FATHER))
                                     (N LEOPOLD))))))

Pragmatic Assumptions:
      I assume by MOZART you mean WOLFGANG-AMADEUS-MOZART.
```

Semantic Case Structure:

```
(CASE   CONSTRUCT-0
        (TENSE          PAST)
        (CONSTRUCTOR        WOLFGANG-AMADEUS-MOZART)
        (CONSTRUCTION   DON-GIOVANNI)
        (RECIPIENT      LEOPOLD (His-0, Dead, Father)))
```

Prototype Structure:

```
(PROTOTYPE   CONSTRUCT
                 (CONSTRUCTOR        ARCHITECT)
                 (CONSTRUCTION   EDIFACE))
```

The domain knowledge with which these prototypical fillers are reconciled must also be considered from the viewpoint of prototypicality and salience. For instance, the metaphor *"Goethe mined the human soul in Faust"* generates three figurative slot:filler juxtapositions, GOETHE AS GEOLOGIST, HUMAN-SOUL AS ROCK and FAUST AS TERRAIN. The first of these should lead a system to generate a metaphoric reconciliation between WRITER (the most salient profession of Goethe) and GEOLOGIST. However, a system that possesses the additional fact that Goethe was also a geologist (after whom the mineral göthite is named) would seem to be at a disadvantage in such a situation, possessing simply *too much* knowledge to recognize the intended metaphoric tension of the utterance. In such cases it seems reasonable to argue the reduced salience of the association between Goethe and geologist is itself a source of tension, prompting the system to still consider the metaphor WRITER AS GEOLOGIST. This view is consonant with that of Ortony (1979), who argues that metaphor is primarily a device for highlighting low-salience attributes in the tenor.

Interestingly, this reconciliation of case fillers with their prototypical expectations is a necessary step even for ostensibly literal uses of a concept. For instance, were the system to be unaware of Mozart's profession, and similarly ignorant of the fact that Don Giovanni is a famous opera, then it would naturally infer Mozart to be an architect, and Don Giovanni to be a building. This form of *semantic promotion* is a valid response given incomplete information, and illustrates that even so-called literal uses of a concept must nevertheless be reconciled with the prototypical expectations as evoked by that concept. The form of reconciliation required in this example is doubtless of a deeper conceptual nature, for the system must create a *theory* of how a composer can be seen to resemble an architect. However, each form of reconciliation (figurative or literal) simply corresponds to a different elaboration process. Since these elaboration processes operate concurrently upon the underlying scaffolding, the *total time constraint* on the interpretation process (see Gerrig, 1989) is never compromised.

Sapper reconciles the current case fillers and expectations as follows:

Prototype Mapping:- Compare W-A-Mozart with Architect.

*If **COMPOSER** is like **ARCHITECT***
Then MUSIC-NOTE is like BRICK
And OPERATIC-ACT is like FLOOR
And DRUM is like PILE-DRIVER
And MUSICIAN is like BUILDER
And MUSIC-RECITAL is like CONSTRUCTION
And OPERA is like EDIFICE
And ORCHESTRA is like CONSTRUCTION-CREW
And LISTENER is like OCCUPANT
And LIBRETTO is like BLUEPRINT

Prototype Mapping: Compare Don-Giovanni with Building.

*If **OPERA** is like **EDIFICE***
Then CHARACTER is like OCCUPANT
And MUSIC-NOTE is like BRICK
And OPERATIC-ACT is like FLOOR
And LIBRETTO is like BLUEPRINT

The lexical definition of "To Construct" specifies a scaffolding structure that is instantiated via co-referencing with the given fillers for the roles AGENT, PATIENT and RECIPIENT (denoted by the indices 1, 2 and 3 in the diagram below). In addition to establishing connections between AGENT and PATIENT, and AGENT and RECIPIENT, the scaffolding also states that the PATIENT is to contain the RECIPIENT. This is a default assumption, since one typically constructs an edifice for another to occupy as a residence. So given the following unification structure for the verb "To Construct":

```
              [ Cat        V
Construct  =  [
              [ Sem        [ Agent      [1] ]
              [            [ Patient    [2] ]
              [            [ Recipient  [3] ]
              [
              [ Scaffold   < connect ( [1] , [2] ),
              [              connect( [1] , [3] ),
              [              cause( [1] , contain( [2] , [3] ))
              [              contain( [2] , [3] )
              [              identity( [2] , Building)
              [              identity( [1] , Architect) >
```

the scaffolding structure of CONSTRUCT-0 can be instantiated as follows.

Conceptual Scaffolding of CONSTRUCT-0 :-

```
<<CONNECT     Wolfgang-Amadeus-Mozart   Don-Giovanni>>
<<CONNECT     Wolfgang-Amadeus-Mozart   Leopold>>
<<CAUSE       Wolfgang-Amadeus-Mozart
                   (CONTAIN     Don-Giovanni    Leopold)>>
<<CONTAIN     Don-Giovanni    Leopold>>
<<IDENTITY    Leopold              Dead-Father>>
<<IDENTITY    Don-Giovanni    Building>>
<<IDENTITY    Wolfgang-Amadeus-Mozart      Architect>>
```

Identity constructs allow substitutions to be made in other parts of the scaffolding, thus allowing the system to construct, among others, the following:

```
<<CONTAIN     Building          Dead-Father>>
```

In effect, identity constructs allow an early form of structural elaboration to take place, fleshing out aspects of the utterance that are implicit in its surface form. This extra structure can in turn provide the grist for deeper conceptual elaboration. In this case, Scout can be evoked to find a structural analog of a building that contains a dead father, determining a concept such as MAUSOLEUM or SHRINE to be a good match. Thus the system can infer the additional scaffolding element:

```
<<IDENTITY    Building          Mausoleum >>
```

Since each new identity construct allows further substitution to occur, the scaffolding element is produced:

<<IDENTITY Don-Giovanni Mausoleum >>

This identification of Mozart's opera with a mausoleum may cause comprehenders to reconsider Don Giovanni as an altogether darker and foreboding affair, by *promoting* its associations with DEATH, DECAY, RETRIBUTION and the AFTERLIFE accordingly (in the vein of Ortony, 1979).

Domain knowledge can now be used to elaborate the scaffolding with specific semantic relations that indicate how its spatial connections and containments are to be conceptually realized. We picture this as a relatively straight-forward task, since the structure of conceptual memory will readily reveal the most specific relationship that is applicable in each case. For instance, memory will record the fact that composers write operas via the link COMPOSER—*create*→OPERA, and since Mozart is known to be a composer, the system can elaborate the link between Mozart and Don-Giovanni with the relation COMPOSER. The scaffolding thus becomes labeled as follows:

Scaffolding Interpretation / Elaboration :-

CONNECT:: W-A-Mozart →O← Don-Giovanni as Composer.
CONNECT:: W-A-Mozart →O← Mausoleum as Architect.
CONNECT:: Leopold →O← W-A-Mozart as Father.
CONTAIN:: Leopold →☐ Don-Giovanni as Character.
CONTAIN:: Leopold →☐ Mausoleum as Occupant.

In general, the system attempts to use those relationship labels that are explicitly stated in the utterance. So, while two people such as Leopold and Mozart can conceivably relate in many different fashions (such as MENTOR, MANAGER, FATHER, FRIEND, COLLEAGUE, etc.), the system operates here with that relation explicitly stated in the utterance, FATHER.

Of more interest is the connection between LEOPOLD and DON-GIOVANNI. Since it is known that Mozart is the composer of this opera, the system must satisfy itself by linking Leopold to the opera not as COMPOSER (again, the valence for this role is one) but as a character, and thus to the mausoleum as an occupant. Indeed, some knowledge of the characters of the opera should allow the system to determine a more specific character mapping for LEOPOLD. Employing Scout with the character list of the play as a retrieval filter, the system should associate LEOPOLD with the character of OLD-GENERAL, since both are STERN, BITTER, OLD, and DEAD and each possesses a wayward son. This association can then be further analysed to find additional metaphoric import, by reinvoking Sapper on the combination LEOPOLD and GENERAL. Since Leopold was himself a well-known composer, the salience of this fact may in turn motivate the metaphor COMPOSER AS GENERAL, producing the mappings of Figures 4,5,7 and 8.

Theoretical and Empirical Issues

A complexity analysis of Sapper is reported in Veale *et al.* (1996a,b; 1997), which reveals Sapper to be a well-behaved polynomial-time algorithm. Veale *et al.* (1995,1996b, 1997) also compare Sapper with the SME model of Falkenhainer, Forbus and Gentner (1989), and with the ACME model of Holyoak and Thagard (1989). The results of our experiments—which involve over one hundred metaphors of the COMPOSER AS GENERAL variety in the profession domain—reveal the latter two models to under-perform considerably when dealing with object-centered metaphors of a *noun:noun* variety. We note that SME typically performs excellently on a wide-range of other metaphor forms, such as those that involve causally-rich narratives. But since *noun:noun* metaphors are as prevalent, if not more so, as event-centered *verb:verb* metaphors, we suggest that competence on such metaphors is an necessary element of any integrated model of metaphor in a linguistic context. Further support is given to this contention in a later section.

The Tractability of Sapper and Scout

It may seem that Sapper and Scout are attempting to reconcile two antagonistic positions at once, in claiming to perform a thorough and structured search of long-term semantic memory while also claiming that such a process is computationally tractable. The key to resolving this apparent contradiction lies in how Sapper and Scout organize semantic memory. Both algorithms are *memory-situated* in the sense that they do not approach memory from *outside*, as external processes wishing to manipulate memory as a separate and distinct database of facts; rather they work from *within* memory, at fixed positions as given by the tenor and vehicle concept nodes of a particular metaphor. So rather than perform an exhaustive search of all the structural possibilities inherent in a given configuration of memory—an immense task to be sure (but one that the ARCS system nonetheless commits itself to)—Sapper and Scout need only search the semantic neighborhood in which they are initiated.

As illustrated in Figure 6, Sapper uses spreading activation to search only that semantic vicinity within a horizon *H* linkages from the tenor matriarch. Denoting the average branching factor or arboricity of a memory node in this horizon as B, the complexity of this search is thus $O(B^H)$. If a bridge or dormant connection to the vehicle domain is found within this horizon, Sapper then attempts to unfold the original tenor-domain pathway into the vehicle domain from this connection, in an attempt to reach the vehicle matriarch. This unraveling process has a worst case complexity of $O(B^H)$, but on average we can expect the unfolding process to be much less expensive than this, as most tenor-domain pathways will not be mirrored completely in the vehicle domain. If we denote as *k* the probability that a particular tenor pathway actually meets a bridge or dormant linkage into the vehicle domain within the horizon *H* , where $k << 1$, then the complexity of Sapper's path-finding stage is at most $O(kB^{2H})$. Prior spreading activation in the vehicle domain serves to keep the value of *k* low, since only bridges leading to a marked vehicle domain concept will be pursued. Scout does not exploit this additional constraint, so we can expect *k* to be closer to 1 on average for this algorithm than for Sapper. Thus, Scout's complexity is closer to $O(B^{2H})$.

The Sapper seeding process is a straightforward sorting task of complexity $O(N\log_2 N)$, where *N* is the number of partial interpretations found during the path-finding process. Once sorted, the final merge stage is a linear one, of complexity $O(N)$, indicating the overall Sapper algorithm to have a complexity of $O(kB^{2H} + N\log_2 N + N)$. Since *H* is a predetermined system constant ($H = 6$ allows for metaphor interpretations containing six levels of nesting), Sapper's complexity is thus a polynomial in *B*, a variable which reflects the average density of memory inter-connection in the tenor domain. While a setting of $H = 6$ may seem somewhat limiting if one views memory as a geometric landscape, one can nevertheless become very removed from a given conceptual starting point by following a chain of six semantic relations in any direction. For instance, Umberto Eco's highly creative and mapping-rich metaphor OPERATING-SYSTEM AS RELIGION (see Eco, 1994) requires just five levels of recursion when processed by Sapper. In general, we have found that H = 6 has been adequate for even the most striking and apparently abstruse metaphors that Sapper has had to process. In this respect then we do not consider it hyperbolae to suggest that Sapper and Scout can search the depths of semantic memory while nevertheless remaining within tractable computational bounds.

The Reality of Conceptual Bridges

Empirically, Sapper's bridge building aspect is consistent with several findings in psychology. Whereas many metaphor models exploit existing associations between concepts (e.g., Fass, 1988; Way, 1990; Martin, 1991), the integrated model *creates* new associations amongst concepts (through Sapper's bridge building mechanism). Camac and Glucksberg (1984) found evidence for the latter position using a word-recognition paradigm. Subjects were presented with several different types of materials to recognize: (i) random letter strings (non-words); (ii) random non-associated word pairs (e.g., SURGEONS : Jails) (iii) word pairs with an a priori association (e.g., DOCTORS: NURSES); (iv) word pairs with a metaphoric association (e.g., Jobs: Jails); and (v) random non-metaphoric non-associated pairs (e.g., DOCTORS: SPOONS). As has been found before in many reaction time studies, word pairings with *a priori* associations were recognized as valid English, being reliably faster than non-associated pairs and non-word pairs. However, no reliable reaction-time difference was found between metaphoric and non-metaphoric word-pairs, strongly suggesting that metaphor word pairs must not have been *a priori* associated in the minds of the test subjects, but created by the very act of interpretation. Kelly and Keil (1987) provide other supporting evidence for this view by showing that inter-concept similarities change after those concepts have been involved in a metaphoric interpretation. The compelling nature of bridging is also reinforced by its appearance in different guises in several different models (see Hofstadter *et al.*, 1995; Hummel and Holyoak, 1996).

But, perhaps the strongest psychological claim we would make for the idea of bridging is one that does not have strong experimental support, but is intuitively compelling; namely, that metaphors and analogies are frequently used as the basis of persuasive arguments that literally *change people's minds* (see Veale and Keane, 1994, Veale, 1998). The anecdotal support for such a stance is the weight given by politicians to the use of the right analogy (usually in wartime) to establish support for their cause. By changing the physical structure of the connections between representations, the integrated model explicitly models this mind-changing capability. Indeed, it is this same mechanism that distinguishes fresh from tired metaphors. For without bridge-building and the representational traces such a approach to metaphor will necessarily lay down in memory, all metaphors would seem equally fresh. These traces of previous interpretations allow a system to recognize established metaphors (i.e., existing bridges) from the more novel tropes that *recruit* and build new bridges upon these established foundations (see Lakoff and Johnson, 1980; Martin, 1991; Veale and Keane, 1992; Fauconnier and Turner, 1998).

Is bridging too profligate?

Though it may seem clear that metaphors should leave some representational trace of themselves in memory, it is not at all obvious that the notion of a bridge advanced in this paper has any cognitive reality in this respect. For this reason we propose the simplest possible representation for a conceptual bridge, while at the same time ensuring that the structures we do provide (i.e., triangulation and squaring) are in some sense *canonical*. For instance, it may be that the structural criteria for laying new inter-domain connections are considerably more complex, and thus more discerning, than triangulation and squaring, resulting in the creation of fewer bridges. Whatever these criteria, such complex constructors can nonetheless be seen as compositional uses of the triangulation and squaring rules, since any structurally consistent lattice structure can be built from these rules. So as we learn more about the conditions under which domains become linked in memory, Sapper too can evolve.

Another potential concern with the notion of bridging as we have advanced it is the profligacy of the triangulation and squaring rules. At first blush it may seem that these rules should lay dormant bridges between all pairs of concepts, leading to an intractable explosion in the size of semantic memory. In practice however, this concern has not actually been realized (but even if it were, the effect would be polynomial over the size of the memory, and not exponential, since a memory of N nodes contains just $O(N^2)$ possible pairings). For instance, in a semantic memory representation of fifteen profession concepts (such as SURGEON, BUTCHER, etc.) comprising 300+ concept nodes and 1600+ semantic relations, automatic triangulation and squaring added on the order of 2300 dormant bridges, far less that the 45,000 predicted by a worst case analysis (see Veale, O'Donoghue and Keane, 1996).

We note also that Sapper's use of dormant bridge structures is a philosophical choice that does not particularly effect the performance or competence of the algorithms discussed in this paper. As described in Veale and Keane (1998), we adopt the view that metaphors are conceptual structures nascent in memory before they are actualized, so that memory itself actively cooperates in the recognition of metaphor. But like all space versus time tradeoffs, the onus of potential bridge recognition can be shifted from memory to the Sapper algorithm itself (if empirical evidence suggests that it should), saving space in memory while adding to the performance time exhibited by the mapping process. In the absence of pre-compiled bridging, the triangulation rule can instead be *fired* at every concept newly visited by the Sapper algorithm. In this way, transient dormant bridges will be laid down in the context of specific mapping problems, and can be removed after an interpretation is established (dormant bridges would thus act as a mapping-level scaffolding structure). Since Sapper visits a polynomial number of nodes (relative to the size of the problem), the triangulation rule would only fire a polynomial number of times. And since it would be constrained to consider only those vehicle domain nodes previously visited by the spreading activation process, each triangulation would consider a polynomial number of potential bridgings. Clearly then, Sapper would still maintain a polynomial complexity if the space/time tradeoff were abolished. Indeed, we have tested Sapper on all of our metaphor examples in this manner, and have experienced no noticeable drop in performance, while competence is preserved entirely.

Empirical Observations Arising from Lexico-Conceptual Integration

Interestingly, this integrated approach conforms to important empirical constraints that have been placed on the metaphor interpretation process by experimental research in psycholinguistics (see Hoffman and Kemper, 1987, for a review of the major constraints, as determined by reaction time experiments). Firstly, because all language utterances—whether ostensibly literal or openly metaphoric—must undergo both a scaffolding and an elaboration stage of processing, the approach does not posit an additional processing stage for metaphors. Longer interpretation times for metaphoric statements as compared to their literal equivalents is thus neither a logical prediction or an empirical outcome of the model. In this respect the approach clearly conforms to the *Total Time Constraint* as discussed in Gerrig (1989).

Secondly, it has been empirically observed that metaphors which are open to significant elaboration over time are nevertheless often comprehended in very time-limited contexts, such as cinema, theatre and even in everyday conversation. Clearly then our metaphor faculty is one which can be engaged at different levels of analysis, producing minimal interpretations at a moment's notice, or richly elaborated interpretations in extended time contexts (see again Gerrig, 1989). Indeed, significant time lapses have been observed between the initial, immediate comprehension of a complex metaphor and its fuller appreciation at a later stage. The two-tier organization of a scaffolding structure explains such differences in time and richness: simply, a scaffolding provides an immediate first cut interpretation of a metaphor, remaining in long-term memory until it can gradually be elaborated and replaced over time.

Finally, there is also convincing empirical evidence for the proposition that metaphor not only exploits existing relations between concepts, but actually invents such relations itself when needed (see for instance Camac and Glucksberg, 1984). In line with this observation, the conceptual scaffolding model has been designed to allow not only for situations where existing domain knowledge is used to elaborate a constructor (for instance, as when the collocation MICROSOFT→O←IBM is elaborated using the relation PARTNERSHIP), but also those situations where two concepts are linked by a constructor but no suitable conceptual relation seems to apply (e.g., BULGARIA→O←PUPPET in the phrase *"Bulgarian puppet"*). In these situations Sapper can generate a mapping theory of how those concepts relate at a finer resolution of knowledge, e.g., how the concepts BULGARIA and PUPPET relate to each other in terms of Bulgaria's fraught relationship with Russia. This issue is discussed in depth in the following sections.

Advantages Arising from A Unified Model

A unified system is not restricted to treating metaphor as a wholly lexical or wholly conceptual phenomenon, but for each metaphoric utterance can craft an interpretation that best exploits its available knowledge concerning the elements of that utterance. For instance, consider the metaphor *"Cold fusion is not kosher physics"*. Now, a system might define the concept KOSHER at multiple levels of detail: at the lexical level, as an informal gloss for LEGAL and PROPER (i.e., something that "checks out"), and at the conceptual level, in terms of specific Judaic dietary laws and religious customs. So depending on how much the system knows about the concept PHYSICS, either a projective gloss or a deeper mapping-based interpretation (one that elaborates the conventional metaphor IDEAS ARE FOOD) can be generated. However, for an utterance like *"Cold fusion is kosher pork"*, only the conceptual level can produce a sensible interpretation, as otherwise the oxymoronic collocation of KOSHER and PORK would not generate the appropriate contradiction (i.e., that the very idea of cold fusion is deeply flawed).

In this section we briefly consider three problematic aspects of metaphor interpretation that we believe are similarly best understood via the organized interplay of lexical semantics and conceptual knowledge.

Literal pretence

Barnden (1998) has noted that even metaphors which seem to marshal whole systems of metaphoric objects (such as "Mary *overhauled* her theory from the *spark plugs* to the *oil-filters*") do not necessarily require the system to find tenor domain counterparts for these objects. Instead, inferences can be made in the vehicle domain (e.g., that Mary was *thorough* in her overhaul) and transferred to the tenor without knowing precisely what SPARKPLUG and OILFILTER correspond to in the THEORY domain. Though a system should recognize (via structure-mapping, we would argue) that the utterance is an elaboration of the FORMALTHEORY AS MECHANISM metaphor, the system should not rely entirely on domain mappings as the basis of an interpretation.

It seems then that a metaphor will sometimes construct a scenario that has neither a literal nor an analogical interpretation. In such cases, it is futile to seek a meaning directly in the source domain, or to establish, by analogical mapping, a set of target counterparts and seek a meaning in the target domain. For instance, consider the metaphor *"Hitchcock's mothers have the psyches of Sherman tanks"* (an elaboration of one considered earlier). Tanks do not literally have psyches, and one would be hard-pressed to determine the mechanical counterpart of one, yet the metaphor is clearly meaningful. The idea of PSYCHE is here used as a domain filter, to direct the listener to focus on those psychological aspects of MOTHER that might also, by means of domain incongruence, be seen as applicable to TANK. The metaphor is *not* intended to convey the cliched view of mothers that a vehicle like BATTLESHIP would readily convey, that of large, lumbering steely gray opponents.

Following Barnden (1998), it seems that a system must engage in a *literal pretence* that tanks do actually possess psyches. We would argue that since scaffolding structures are pre-interpretative, they can serve as non-committal representations for this kind of pretence. Thus, a system can posit a common psychological ground between MOTHER and TANK via the scaffolding MOTHER□← PSYCHE →□TANK, and thus achieve a mapping between both. In effect, the scaffolding temporarily creates a basis for triangulation by positing PSYCHE as a common associate, so that given an appropriate projective strategy, the system can then transfer domain-incongruent descriptions of a psychological nature, such as DETERMINED, OBSTINATE, DRIVEN and AGGRESSIVE, from the domain of TANK to that of MOTHER.

Non-compositionality of interpretation

As the previous example illustrates, metaphoric utterances are not always strictly compositional, inasmuch as the global meaning of the whole cannot be constructed piecewise from the local meaning of the parts. The phrase *"the psyche of a Sherman tank"* has no coherent literal or analogical meaning as a stand-alone term, and can only be interpreted in the context of the tenor it is used to describe. From a scaffolding perspective, this means that a scaffolding structure cannot likewise be interpreted as a composition of the interpretation of each individual scaffolding operator. Thus, an elaboration strategy must process the scaffolding MOTHER□←PSYCHE→□TANK as a whole, for the meaning of MOTHER□←PSYCHE is simply banal (mothers have psyches) while the meaning of PSYCHE→□TANK is too diffuse to be considered well-defined (tanks do not have psyches, or anything remotely corresponding to psyches). Only when taken together do these individual structures yield any real meaning. Ironically then, while a scaffolding structure is *constructed* compositionally (from lexical semantics), it cannot be *interpreted* compositionally.

Compositionality of metaphoric meaning is a key claim of models of metaphor such as that of Lytinen *et al.* (1992), who argue that the literal meaning of an utterance is constructed compositionally at each level of syntax (i.e., the meaning of a verb-phrase is composed from the meaning already assigned to its subordinate noun-phrases), and that any metaphoric meaning is then constructed from this meaning via mapping rules. However, as we have argued, a system may not be able to ascribe metaphoric meaning to individual noun-phrases without first looking at the full context in which those phrases occur. For instance, the vehicle phrase *"kosher pork"* is meaningless in itself, but can be seen as meaningful in a larger context when the issue is the purported meaninglessness of the tenor.

In particular, the compositionality view is vulnerable to utterances in which *modifier-switching* occurs. For instance, in sentences like *"Viagra is sexual rocket fuel"* and *"Love is metaphysical gravity"*, the vehicle is described with a modifier that semantically belongs to the tenor. Since metaphors establish an identity (or category inclusion) between tenor and vehicle, this switch is meaningful, but only in the global context of the utterance. For locally, one would be hard-pressed to ascribe any specific meaning to a phrase like *"metaphysical gravity"*. This is in fact a secondary metaphor that relies on the primary metaphor (LOVE AS GRAVITY) being interpreted first. So while the compositionality stance conventionally assumes that phrases of the form "A is B C" are interpreted piecewise in the fashion *(A = (B + C))*, many such utterances can only be meaningfully interpreted in the holistic fashion *((A = B) + (A = C))*. The scaffolding structure provides a global context in which such metaphors can be viewed holistically, allowing specific elaboration strategies to work either at the level of individual operators (e.g., MARY □← INFLUENZA) or at the level of compound operator structures (e.g., ELEPHANT→O←TANK→□ DESERT)

Coherence among counterparts

A conceptual scaffolding ties the key elements of comprehension together, allowing elaboration processes to work in unison to achieve a coherent overall interpretation for a metaphor. Thus, when various elements of a metaphoric utterance suggest potential analogical remindings, coherence can be enforced between these remindings.

Consider the utterance *"Hannibal's tanks attacked his Monty in the battle of Zama"*. This metaphor primarily establishes an analogy between TANK and ELEPHANT, and between MONTGOMERY and SCIPIO AFRICANUS (Hannibal's ultimate vanquisher at Zama). The elaboration strategy for interpreting *"Hannibal's tanks"* (TANK→□HANNIBAL) and *"Hannibal's* [his] *Monty"* (MONTGOMERY→□HANNIBAL) is the *domains as containers* strategy described earlier. Scout's use of spreading activation in this strategy highlights a number of potential analogs for Hannibal himself: if elephants can be seen as tanks, then Hannibal can be seen as a famous tank commander such as Rommel, Montgomery or Patten. Likewise, the battle of Zama can be seen as the battle of Alemain, the battle in which Montgomery defeated Rommel.

But MONTGOMERY is already governed by HANNIBAL in the phrase *"Hannibal's Monty"*. Thus, Scout must find an analog for MONTGOMERY in the domain of HANNIBAL. Since the BRITISHEMPIRE is a good analog for the ROMANEMPIRE, SCIPIO will be seen by Scout as a rich analog for MONTGOMERY. If MONTGOMERY maps to SCIPIO, then—by the principle of isomorphism underlying metaphor coherence—HANNIBAL cannot also map to MONTGOMERY. But as ROMMEL is a highly salient opponent of MONTGOMERY, just as HANNIBAL is of SCIPIO, Scout should prefer ROMMEL as a mapping for HANNIBAL rather than PATTON (or any other general for that matter). This in turn should involve a subordinate mapping of ZAMA to ALAMEIN (both are, respectively, the North African sites of decisive battles in the careers of Hannibal/Rommel and Scipio/Montgomery).

The motivations for these mappings arise from the lexical level of analysis (e.g., note how the genitive construction lays down a containment operator between TANK and HANNIBAL), but the mappings can only be resolved coherently at the conceptual level. In a unified model then, the scaffolding acts as a global workspace for coordinating different mapping processes (i.e., different Sapper and Scout invocations) while the mapping mechanism itself ensures the coherence of different analogical correspondences by enforcing the basic constraint of graph isomorphism.

How Metaphor Informs Syntax

A basic tenet in the integration of linguistic and knowledge-based models of language is that high level processes such as syntactic analysis must not only shape the processes that occur at the semantic and conceptual level, they must respond to them accordingly if structural and referential ambiguities are to be resolved. This may seem a trite statement of established fact, as many current models of language—such Pollard and Sag's HPSG (1987)—place syntax and semantics on an equal footing to allow this interplay to occur (while remaining neutral on issues of process ordering and interpretation strategy), but this parity is often absent in models of metaphor interpretation. For instance, while the model of Aarts and Calbert may sometimes disallow a certain adjective:noun combination on the basis of marker incompatibility or non-extensibility, their model makes no allowance for the possibility that the semantic/conceptual level may sometimes not go so far as to strongly disallow, but merely not prefer, a particular combination. That is, ambiguities may not necessarily be resolvable on the basis of all or nothing information, but on the basis of one conceptual structure being more preferable than another.

Indeed, the case made by the conceptual level may sometimes be so compelling as to cause the syntactic analysis to accept what seems to be a blatantly ungrammatical analysis. For instance, Pollard and Sag (1987) cite the example sentence "*The hash browns wants his bill*" in which a metonymy is so transparent as to persuade the syntax component into accepting an analysis where number agreement is violated. The case made by the conceptual analysis is seemingly strong enough for us to accept the phrase "*The hash browns*" on an almost idiomatic basis as referring directly to a person. This is a situation where the very grammatical features of syntax are plastic in the face of conceptual analysis, but other less drastic situations abound.

Metaphoric Ambiguities

When the conceptual criterion on which a syntactic ambiguity turns is metaphoric in nature, this situation is exacerbated. This is for the most part due to the separate progression of the strong linguistic and strong conceptual approaches to metaphoric language. For instance, consider a metaphor analyzed in Way (1991), "*Bulgaria is a Russian puppet government*". Way's system resolves this metaphor by connecting the concepts BULGARIA and PUPPET via a common supertype *ENTITY-OR-THING-WHICH-IS-CONTROLLED-FROM-ABOVE*. This solution is reminiscent of Russell's MAP, which too seeks to resolve metaphors by reducing the breadth of world knowledge that surrounds a concept such as PUPPET to a lexico-semantic abstraction (e.g., compare Russell's definition of TORPEDO with Way's definition of PUPPET). But note how the structural ambiguity of the utterance, observable at a basic linguistic level, has been side-stepped. In the compositional analysis shown below, the noun phrase "Russian puppet government" actually supports two different analyses: (i) projects the features of "puppet government" as its head, while (ii) projects only those of "government".

(i) $[_{NP}$ Russian $[_{NP}$ puppet government]]

(ii) $[_{NP}$ $[_{NP}$ Russian puppet] government]

The second analysis, in positing the well-formed sub-construct "Russian puppet", is more likely than the first to prompt one to think of Russia as a puppet and thus to construe the larger phrase as referring to a Moscow government ruled from outside. The first analysis, however, associates "puppet" directly with "government", not "Russia"; we are thus less likely to consider Russia as a puppet, and more likely to construe the larger phrase as referring to a government ruled from the outside by Moscow. (We make no strong claims for these specific interpretations, as they are most likely listener-dependent, causing some native speakers to hold opposing views; what we do claim, however, is that each analysis generates a different interpretation, and that this resultant ambiguity is resolvable only at the knowledge-level).

A major factor in our preference for the former analysis is our contingent world knowledge of eastern European politics. Russia—a behemoth in economic and military terms that until recently controlled almost all of eastern Europe—does not sit at all well in the role of puppet: simply put, if Russia is a puppet, who could possibly assume the role of puppet master? In contrast, due to the immensely smaller size and economy of Bulgaria, and our knowledge that Russia does indeed (or once did) exert military and economic pressures upon it, the juxtaposition of PUPPET and BULGARIA actually makes for a more consonant metaphor interpretation.

Of course, there are also semantic rather than conceptual reasons for preferring the former analysis: if we assume that "puppet" refers to Russia, then the two become co-referential and Russia shares the role of semantic head of the noun phrase. The effect of this is to equate the Russian government with Bulgaria, producing the anomalous claim that Russia and Bulgaria are one and the same. Nevertheless, the problem also recurs at a sentential level in the utterance "*America withdrew its ambassador from the Russian puppet government of Bulgaria*". In this case there are no syntactic or semantic cues to resolve the ambiguity globally, so two alternative syntactic analyses are produced. The conceptual scaffoldings extracted from the alternate analyses are illustrated in Figure 11.

Figure 11 around Here, "Russian/Bulgarian Scaffoldings"

The ambiguity is essentially one of deep case assignment: what assignment of the concepts RUSSIA and BULGARIA to the conceptual roles PUPPET and PUPPET-MASTER best resonates with the contents of long-term memory? The relevant contents of memory as relating to this ambiguity are illustrated in Figure 12, demonstrating very compelling grounds for viewing Russia as puppet master rather than simple puppet.

Figure 12 around Here, "Russian/Bulgarian Memory Pattern"

Metaphor and Case Grammar

As can be seen from the memory fragment of Figure 12, and the resulting Sapper analysis of Figure 13, a rich metaphoric mapping emerges from the juxtaposition of the concepts RUSSIA and PUPPET-MASTER. Note how Sapper incorporates knowledge of different types, not simply political, to build its mapping theory. For instance, a speaker (or listener) with a representation as baroque as that of Figure 12 may consider that there are geographical grounds for linking the concepts BALKANS and BALCONY, as both are high, imposing structures. The speaker may also notice phonetic grounds for this mapping, as both share a similar pronunciation for their first two syllables, "bawl" + "kan". Speakers frequently combine such varied knowledge sources, both conceptual and perceptual, as evidenced by the popularity of punning humour. In this context, one might even say that *"NATO observers in northern Greece have a Balkan seat at the puppet show*". Not every agent will make such connections, of course (indeed, many would find such an interpretation tenuous at best). The Sapper model simply claims that if the basis for such connections is present, and within Sapper's horizon setting, they will be opportunistically recognized and included into the overall mapping by Sapper's spreading activation process.

Indeed, we advance this as a major advantage of the Sapper model of conceptual metaphor: while bridges might be explicitly inferred by the triangulation rule on the basis of shared conceptual similarities between domains, we also envisage an active role for lower-level, sub-symbolic perceptual processes in the creation of these bridges. But the ability of the squaring rule to extrapolate these perceptual similarities into higher conceptual juxtapositions means we do not have to necessarily concern ourselves with the mechanics of these lower-levels. Overall then, Sapper acknowledges the figurative effects of the interaction that occurs between perceptual and conceptual levels of cognition, as elucidated by metaphor researchers such as Beck (1976), Harnad (1982) and Indurkhya (1992), while conveniently allowing our research to proceed wholly in the conceptual realm.

Interestingly, the mapping theory generated for RUSSIA and PUPPET-MASTER subsumes that of the metaphor BULGARIA: PUPPET, while also being demonstrably richer than those produced for the juxtapositions BULGARIA: PUPPET-MASTER and RUSSIA: PUPPET. It is on this basis, then, that we argue that a conceptual model such as Sapper can act as an arbiter of case assignment in language comprehension systems.

Figure 13 around Here, "Russian/Bulgarian Mappings"

This posited role for metaphor runs counter to conventional wisdom in the use of case grammars. Wilks (1975,1976,1978), for instance, introduces the notion of a *preference semantics* which models the semantics of deep cases in terms of soft preferences rather than hard constraints. The central idea here is that since metaphors are likely to stretch the conventional uses of deep cases, we should make the semantic criteria governing case assignment more flexible and thus more robust. For instance, the agent case of the verb "To Drink" may prefer an animate entity, but will accept a non-animate filler if presented with one in a metaphor such as "*My car drinks gasoline*". Likewise, if the verb "To Choose" also prefers a sentient agent, it too will accept a non-sentient filler if presented with no other choice, as in "*The hurricane chose a Monday to hit town*". The traditional view thus views case grammars as catering to the literal uses of words, but accommodating metaphors whenever they arise. The revised view that we advocate here turns this scenario of its head: metaphor should be seen as a core (rather than peripheral) phenomenon that mediates between a case role and its filler.

This alternate view gives metaphor an explanatory role in the process of case filling, in contrast to a traditionally deviant role that case filling processes must *explain away*. This explanatory role in turn leads to a theory that can more truly claim to understand why certain concepts can fill certain case slots and not others. For instance, HURRICANE can fill a case normally occupied by an instance of PERSON because one can construct a convincing mapping theory to reconcile the concepts of HURRICANE and PERSON. This mapping theory would make apparent certain shared features of both, such as unpredictability, mobility, place-of-origin/birth-place, and identification with personal names (it helps that hurricanes are frequently personified). Likewise, a ROBOT can fill a case which expresses a preference for HUMAN fillers, via the very rich mapping theory that underlies the ROBOT AS HUMAN metaphor.

Figure 14 illustrates the relevant elements of long-term memory that are called into play when the verb "To Open" is used in a decryption context, as in "*Bob opened his encrypted files with a password*". Assuming that the deep case structure for "To Open" is specified relative to some prototypical usage—such as opening a room/door with a key—then case analysis of this utterance relies upon the system being able to reconcile the concepts ROOM and ENCRYPTED-FILE, and PASSWORD and KEY.

Figure 14 around Here, “Code / Lock Metaphor”

Competing case-structure analyses of a syntactic construct (either a noun phrase or sentence) should be evaluated then relative to the *structural fit* which they yield with the contents of long-term memory. This fit is best evaluated using a knowledge-based model of metaphor like Sapper, one with the analogical power to construct mapping theories between the domain of a given case filler and the domain of the preferred filler, to explain precisely why the given concept can act as a suitable filler. This is somewhat similar to Metallel's use of converging paths to reconcile fillers to their role expectations. However, rather than use single paths, Sapper uses structural mappings—whole systems of consistent paths—to perform this reconciliation. In effect, Sapper acts as a *category inclusion* mechanism (in the sense of Glucksberg's theory) to show how the filler concept can be included under the category of the role expectation. The richness of a mapping theory, e.g., the number of mappings it contains, thus provides an excellent *‘credibility measure’* for case assignment.

<u>Deeper Understanding in Case Grammars: A Proposal</u>

In clarifying the nature of the case filling mechanism by introducing metaphor as a core process, we can in turn clarify the role of case grammar in text comprehension, and explain just how case-grammar interpretations can be said to *understand a text*. Metaphor systems which represent meanings in terms of deep case structures can be positioned along a continuum defined between two extremes: at one extreme we can characterize a *laissez-faire* attitude toward metaphoric meanings, in which the meaning of a metaphoric utterance resides in those case-frames which are directly indicated by the lexical semantics of the utterance, regardless of how many preference-violations this may entail. For instance, in a laissez-faire system the meaning of “*my car drinks gasoline*” is represented in terms of a DRINK case-frame rather than a more accommodating CONSUME case-frame. At the other end of the continuum we can characterize a *repair* attitude which states that whenever a metaphoric meaning violates one or more case preferences, the system should attempt to remap the meaning into a set of case-frames for which no preferences are violated. Thus, in the thirsty car example, a repair system would consider a CONSUME frame to more accurately represent the meaning of the utterance. McKevitt

(1991) may be the only author to seriously advocate the use of the *laissez faire* approach, while the repair approach is characteristic of many models, such as those of Wilks (1978), Fass (1988), Lytinen *et al.* (1992), and, to a lesser extent, Iverson and Helmreich (1992).

But both of these perspectives miss the mark. The laissez-faire approach is simply too *representationally inert* to ever make any real claim to understanding a text. Likewise, the repair approach seems to preclude metaphoric meanings from ever being represented as true meanings at the knowledge-level of a system, instead necessitating that their figurative elements be spirited away (i.e., repaired). The approach we describe in this paper is positioned midway between both these extremes. As with the laissez-faire approach, we advocate that the case-frames directly suggested by the metaphoric usage of a word, rather than some literal replacement (e.g., DRINK versus CONSUME), form the ultimate representation of a metaphoric utterance. However, we eschew the innertness of the laissez-faire approach, advocating instead that preference violations are explained via analogical mapping theories which reveal the non-obvious, underlying similarities between a preference and its assigned filler. In this view a car *can* drink gasoline, precisely because like organic creatures it too has a digestive system—albeit a mechanical one—that includes a mouth (GASCAP), throat (GASLINE), stomach (GASTANK) and bowel (SUMP). In the end then, though we distance ourselves from his *anomaly-driven* view of metaphor, Wilks (1978) still has the last word: only by making preferences more active can a system hope to adequately handle a phenomenon as pervasive and slippery as metaphor.

Summary and Conclusions

We can summarize our earlier Mozart example, and what it tells us about the requirements of metaphor processing, by noting that the mapping and interpretative processes underlying Salieri's observations clearly require multiple perspectives on the structure of semantic memory. The Conceptual Scaffolding model provides a *rarefied* top-down view in which high-level spatial schemas are imposed on the utterance to organize it into a pattern of conceptual associations, disassociations, containments and other spatial constructs, while elaboration strategies bring mechanisms such as Sapper and Scout to bear on each such association in a bottom-up fashion.

We believe that the notion of an intermediate scaffolding structure is intrinsic to the success of a unified model. Because a scaffolding is pre-interpretative, it can be constructed from lexical information yet still convey—through spatial metaphor—a broad picture of the utterance's meaning as a whole. This picture can then inform the conceptual processes that flesh out the deep meaning of the utterance, so that processes working on one part of the utterance can obtain some perspective on the meaning of the other parts. Without an intermediate meaning level, such processes would be reduced to working in a wholly compositional manner, and as we have demonstrated, this is an inadequate approach for many metaphors.

An integrated approach is necessary not only for the analysis of highly rhetorical expressions (as one might find in a literary work like 'Amadeus'), but of everyday expressions in which the resolution of syntactic ambiguity hinges upon contingent world knowledge. Only in a truly integrated approach can common-sense knowledge about the world, in addition to purely linguistic sensibilities of an abstract nature, be brought to bear on the interpretation process. In many other integrated approaches to language, common-sense has been incorporated into the comprehension process via procedural knowledge sources that work in concert throughout the syntactic and semantic analysis phases (see for instance, the demon-based parsing and understanding systems of Schank, 1975, and Dyer, 1983). Our solution to the problem is, we believe, a more elegant one, inasmuch as it turns what is often conceived as yet another major problem of language analysis, metaphor, into a solution. We have argued how an analogically proficient model of metaphor analysis can be exploited to perform the more general task of *conceptual reconciliation*, whereby one conceptual structure is related to another via a mapping theory which explains (and often *invents*, via the squaring rule) the common ground between both. This notion of reconciliation can be used to test the fit of hypotheses at any level of analysis (e.g., over which noun does an adjective have scope? to which case should a constituent be bound? and so on) against the contents of long term memory. As those contents change (e.g., Russia falls into a crippling civil war, while Bulgaria invades Romania and installs a pro-Sofia '*Bulgarian puppet*' government), so too do the perceived merits of such hypotheses.

Implementation Details

We close by noting the practicality of the approach: the Conceptual Scaffolding architecture has been implemented both as part of the *Twig* concept indexing system (see Veale and Cunningham, 1992) and the *Zardoz* English-to-Sign-Language translation system (see Veale *et al.* 1994, 1996, 1998), each upon a Unix/Common-LISP platform. The word/concept hierarchy of Twig (shared by the Zardoz system) includes over 3000 concept definitions, many of which contain a basic conceptual-scaffolding backbone (e.g., SAFETY is UP; DANGER is DOWN; DIVORCE is a DISCONNECTION; RESCUE is to CAUSE a CONNECTION with SAFETY or another upward state; etc.). This backbone allows Zardoz to infer new gestural codings for words/concepts that have no associated sign in its target sign-language lexicon. For instance, scaffolding operators in the definition of MARRIAGE allow Zardoz to sign the concept MERGER as a *coming together* of companies. In effect, Zardoz signs a corporate merger as a *marriage of companies*, by employing the sign for MARRIAGE with hand-classifiers that designate COMPANY rather than PERSON.

The Sapper and Scout systems too have been implemented, while Sapper has been experimentally compared with SME (the *Structure-Mapping Engine*) and ACME (the *Analogical Constraint Mapping Engine*) in Veale *et al.* (1995;1996a,b;1997), and shown to be a practical and memory-oriented solution to the problem of structure-mapping. All three have been integrated in a single Prolog environment to comprehend the Mozart example considered earlier. Apart from this rather pyrotechnical example, which we cite to illustrate the necessity of integration on an utterance-wide level, we have primarily tested the model upon of sub-sentential structures to explore the more subtle aspects of integration. For instance, we concentrate our primary research energies on an exploration of the role of integrated lexical-conceptual processing in the interpretation of adjective:noun and noun:noun compounds. We have tested around one hundred such examples, of a relatively commonplace but nonetheless complex variety such "war elephant", "browser war", "vegetarian sushi" and "math clinic". We have found the analysis of such compounds to be compelling evidence for the dual role of lexical-semantics and deep conceptual structure-mapping in the understanding of everyday phrases.

A Prolog implementation of Sapper, Scout and the Conceptual Scaffolding models (including the noun compounding elaborative strategies alluded to above) can be obtained from following metaphor-dedicated web-site: *http://www.compapp.dcu.ie/ ~tonyv/metaphor.html*. This site also provides the test-data upon which the comparative tests of Sapper, SME and ACME were performed.

References

Aarts, J. M. and J. P. Calbert. (1979). *Metaphor and Non-Metaphor: The Semantics of Adjective-Noun Combinations.* Tübingen: Niemeyer.

Alshawi, H. (1990). Resolving Quasi-Logical Form. *Computational Linguistics* **16**, 133-144.

Anderson, J. M. (1971). *The Grammar of Case: Toward a Localistic Theory.* London: Cambridge University Press.

Aristotle. (1982). *The Poetics.* Translated, with introduction, by J. Hutton. NY: Norton.

Barnden, J. A. (1992). Taking Commonsense Psychology Seriously. *Computational Intelligence* **8**(3), pp 520-552.

Barnden, J. A. (1998). Concerning the role of analogy in metaphor processing, *in the proceedings of Advances in Analogical Research, eds. Keith Holyoak, Dedre Gentner and Boicho Kokinov. NBU Series in Cognitive Science, Sofia 1998.*

Beck, E. F. (1978). The Metaphor as a Mediator Between Semantic and Analogic Modes of Thought, *Current Anthropology* **19**(1).

Black, M. (1962). *Models and Metaphor: studies in language and philosophy.* Ithaca, NY: Cornell University Press.

Brachman, R. J. (1979). On the Epistemological Status of Semantic Networks, in *Associative Networks: Representation and Use of Knowledge by Computers*, N. V. Findler (ed.). New York: Academic Press.

Camac, K. and S. Glucksberg. (1984). Metaphors do not use associations between concepts, they are used to create them. *Journal of Psycholinguistic Research*, **13**(6).

Charniak, E. (1983). Passing Markers: A Theory of Contextual Influence in Language Comprehension. *Cognitive Science* 7, pp 171 -190.

Collins, A. and E. F. Loftus. (1975). A Spreading-Activation Theory of Semantic Processing. *Psychological Review* **82**, pp 407-428.

Cunningham, P. and T. Veale. (1991). Organizational issues arising from the integration of the Concept Network and Lexicon in a Text Understanding System, in *the Proceedings of the 12th International Joint Conference on Artificial Intelligence.* San Mateo: Morgan Kaufman.

Dyer, M. G. (1983). *In-Depth Understanding.* Cambridge, MA: The MIT Press.

Eco, U. (1984). *Semiotics and the Philosophy of Language.* London: Macmillan Press.

Eco, U. (1994). *La Bustine Di Minerva.* Espresso, September 30, 1994.

Falkenhainer, B., Forbus, K. D., and D. Gentner. (1989). Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence*, **41**, 1-63.

Fass, D. (1988). An Account of Coherence, Semantic Relations, Metonymy, and Lexical Ambiguity Resolution, in *Lexical Ambiguity Resolution: Perspectives from Psycholinguistics, Neuropsychology and Artificial Intelligence.* Small, S., G. Cottrell and M. Tanenhaus, eds. San Mateo, CA: Morgan Kaufmann.

Fauconnier, G. and M. Turner. (1994). *Conceptual projection and middle spaces.* UCSD: Department of Cognitive Science Technical Report 9401.

Fauconnier, G. and M. Turner. (1998). Conceptual Integration Networks. *Cognitive Science*, 22:2. pp 133-187.

K. D. Forbus and D. Oblinger. (1990). Making SME Pragmatic and Greedy, in *the Proceedings of the Twelfth Annual Meeting of the Cognitive Science Society.* Hillsdale, NJ: Lawrence Erlbaum.

Gerrig, R. J. (1989). Empirical constraints on computational theories of metaphor: Comments on Indurkhya. *Cognitive Science* **13**(2), 235-241.

Gibbs, R. J. (1992). Categorization and metaphor understanding. *Psychological Review* **99**, pp 572 - 577.

Glucksberg, S. and B. Keysar. (1990). Understanding metaphoric comparisons: beyond literal similarity. *Psychological Review* **97**, pp 3 - 18.

Glucksberg, S., McGlone, M. S., and D. Manfredi. (1997). Property Attribution in Metaphor Comprehension. *Journal of Memory and Language* **36**, pp 50 - 67.

Harnad, S. (1982). *Metaphor and Mental Duality*, in *Language, Mind and Brain*, T. Simon and R. Scholes, eds. Hilldale, NJ: Erlbaum, pp 189-211.

Hendler, J. A. (1989). Marker Passing over Micro-Features: Toward a Hybrid Symbolic/Connectionist Model, *Cognitive Science* **13**(1).

Hoffman, R. and S. Kemper. (1987). What could reaction-time studies be telling us about metaphor comprehension? *Metaphor and Symbolic Activity* **2**(3), 149-186.

Hofstadter, D. and M. Mitchell. (1988). Conceptual Slippage and Analogy-Making: A report on the CopyCat Project, in *the Proceedings of the 10th Annual Conference of the Cognitive Science Society*, Montréal, Québec.

Hofstadter, D. R. and the Fluid Analogy Research Group (1995). *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. Basic Books, NY.

Holyoak, K. J. and P. Thagard. (1989). Analogical Mapping by Constraint Satisfaction, *Cognitive Science* **13**, 295-355.

Hummel, J. E. and K. J. Holyoak. (1996). LISA: A Computational Model of Analogical Inference and Schema Induction, in *the Proceedings of the Eighteenth Annual Meeting of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum.

Indurkhya, B. (1992). *Metaphor and Cognition: Studies in Cognitive Systems*. Kluwer Academic Publishers, Dordrecht: The Netherlands.

Iverson, E. and S. Helmreich. (1992). Metallel: An Integrated Approach to Non-Literal Phrase Interpretation, Computational Intelligence **8**(3).

Jackendoff, R. (1983). *Semantics and Cognition.* Cambridge, MA: MIT Press.

Jackendoff, R. (1990). *Semantic Structures*. Cambridge, MA: MIT Press.

Johnson, M. (1987). *The Body in the Mind: The bodily basis of meaning, reason and imagination.* Chicago, Illinois: University of Chicago Press.

Katz, J. J. and J. A. Fodor. (1964). *The Structure of Language: Readings in the Philosophy of Language.* Englewood Cliffs, NJ: Prentice-Hall.

Keane, M.T., T. Ledgeway, and S. Duff. (1994). Constraints on analogical mapping: A comparison of three models. *Cognitive Science*, **18**, 287 - 334.

Kelly, M. H. and F. C. Keil. (1987). Metaphor Comprehension and Knowledge of Semantic Domains, in *Metaphor and Symbolic Activity*, **2**(1), p33-51.

Lakoff, G. and M. Johnson. (1980). *Metaphors We Live By.* Chigaco, Illinois: University of Chicago Press.

Langacker, R. (1991). *Concept, Image, and Symbol.* Berlin: Mouton de Gruyter.

Law, K., K. D. Forbus, and D. Gentner. (1994). Simulating Similarity-Based Retrieval: A Comparison of ARCS and MAC/FAC, in *the Proceedings of the Sixteenth Annual Meeting of the Cognitive Science Society,* Atlanta, Georgia. Hillsdale, NJ: Lawrence Erlbaum.

Lyons, J. (1977). *Semantics*. London: Cambridge University Press.

Lytinen, S., R. Burridge & J. Kirtner. (1992). The role of Literal Meaning in the Comprehension of Non-Literal Constructions. *Computational Intelligence* **8**(3), pp 416-432.

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

McGlone, M. S. (1996). Conceptual Metaphors and Figurative Language Interpretation: food for thought? *Journal of Memory and Language* **35**, pp 544 - 565.

McKevitt, P. (1991). An Evolutionary Model of Natural Language Processing, in *the proceedings of AICS'91, the 1991 Irish conference on Artificial Intelligence and* Cognitive Science. Springer Verlag Lecture Notes in Computer Science.

Ortony, A. (1979). The role of similarity in similes and metaphors, in *Metaphor and Thought,* edited by A. Ortony. Cambridge, MA: Cambridge University Press.

Pollard, C. and I. Sag. (1987). *Information-based Syntax and Semantics: Volume 1 Fundamentals.* CSLI Lecture Notes 13, Chicago University Press: Chicago.

Quillian, M. R. (1968). Semantic Memory, in *Semantic Information Processing,* ed. Marvin Minsky. Cambridge, MA: MIT Press.

Reyle, U. (1993). Dealing with ambiguities by underspecification: Construction, representation and deduction. *Journal of Semantics* **10**, 123-179.

Richards, I. A. (1936). *The Philosophy of Rhetoric.* NY: Oxford University Press.

Rosch, E. and C. B. Mervis. (1975). Family Resemblances: Studies in the Internal Structure of Categories, *Cognitive Psychology* **7**, 573-605.

Rosch, E., C. B. Mervis, W. D. Gray, D. M. Johnson and P. Boyes-Braem. (1976). Basic Objects in Natural Categories, Cognitive Psychology **8**.

Russell, S. W. (1976). Computer Understanding of Metaphorically Used Verbs. *American Journal of Computational Linguistics,* 1976. Microfiche **44**.

Russell, S. W. (1992). Metaphoric Coherence: Distinguishing Verbal Metaphor from "Anomaly", *Computational Intelligence* **8**(3), 553-574.

Schank, R. C. (1975). *Conceptual Information Processing.* Amsterdam: North Holland.

Shaffer, P. (1981). *Amadeus*. London: Penguin Books.

Shieber, S. (1986). *An Introduction to Unification-based Approaches to Grammer*. CSLI Lecture Notes **4**.

Tourangeau, R. and Sternberg, R. J. (1981). Aptness in Metaphor, *Cognitive Psychology* **13**.

Veale, T. and A. Conway. (1994). Cross-modal comprehension in Zardoz, an English to Sign Language translation system, in *the Proceedings of the 7th International Workshop on Natural Language Generation*, Maine, July 1994.

Veale, T. and P. Cunningham. (1992). Competitive Hypothesis Resolution in TWIG: A Blackboard-Driven Text-Understanding System, *in the Proceedings of the 10th European Conference on Artificial Intelligence*, Chichester: John Wiley and Sons.

Veale, T. and M. T. Keane. (1992a). Conceptual Scaffolding: Using metaphors to build knowledge structures, in *the Proceedings of ECAI'92, the 10th European Conference on Artificial Intelligence*, Chichester: John Wiley and Sons.

Veale, T. and M. T. Keane. (1992b). Conceptual Scaffolding: A spatially founded meaning representation for metaphor comprehension. *Computational Intelligence* **8**(3), pp 494-519.

Veale, T. and M. T. Keane. (1993). A Connectionist Model of Semantic Memory for Metaphor Interpretation, in *the Proceedings of the 1993 Workshop on Neural Architectures and Distributed AI*, the Center for Neural Engineering, U.S.C. California.

Veale, T. and M. T. Keane. (1994). Belief Modelling, Intentionality and Perlocution in Metaphor Comprehension, *in the Proceedings of the Sixteenth Annual Meeting of the Cognitive Science Society*, Atlanta. Hillsdale, NJ: Lawrence Erlbaum.

Veale, T. and M. T. Keane. (1997). The Competence of Sub-Optimal Structure Mapping on 'Hard' Analogies, in *the Proceedings of IJCAI'97, the International Joint Conference on Artificial Intelligence,* Nagoya, Japan, August 1997.

Veale, T, D O'Donoghue and M. T. Keane. (1995). Epistemological Issues in Metaphor Comprehension: A Comparative Analysis of Three Models, in *the Proceedings of ICLC'95, the conference of The International Cognitive Linguistics Association*, Albuquerque, New Mexico.

Veale, T, D O'Donoghue and M. T. Keane. (1996). Computability as a Limiting Cognitive Constraint: Complexity Concerns in Metaphor Comprehension about which Cognitive Linguists Should be Aware, *in Cognitive Linguistics: Cultural, Psychological and Typological Issues* (forthcoming).

Veale, T, B. Smyth, D. O'Donoghue and M. T. Keane. (1996). Representational Myopia in Cognitive Mapping, in *the Proceedings of the 1996 AAAI workshop on Source of the Power in Cognitive Theories*, Portland, Oregon, August 1996.

Veale, T. and B. Collins. (1996). Space, Schematization and Metaphor in Sign: Sign Language Translation in the Zardoz System in *the Proceedings of AMTA'96, The 2nd conference of the Association for Machine Translation in the Americas,* Montrèal 1996.

Veale, T., A. Conway and B. Collins. (1998). The Challenges of Cross-Modal Translation: English to Sign Language Translation in the ZARDOZ System, to be published in *Machine Translation*.

Veale, T. and M. T. Keane. (1998). Principle Differences in Structure-Mapping, *in the proceedings of Advances in Analogical Research, eds. Keith Holyoak, Dedre Gentner and Boicho Kokinov. NBU Series in Cognitive Science, Sofia 1998.*

Way, E. C. (1991). *Knowledge Representation and Metaphor, Studies in Cognitive Systems*, Kluwer Academic Publishers.

Weber, S. H. (1988). *A Structured Connectionist Approach to Direct Inferences and Figurative Adjective-Noun Combinations*, Technical Report 289, Computer Science department, University of Rochester.

Weiner, J. E. (1984). A knowledge representation approach to understanding metaphors, Computational Linguistics, **10**(1), 1-15.

Wilks, Y. (1975). A preferential, pattern-Seeking, semantics for natural language inference, *Artificial Intelligence* **6**, 53-74.

Wilks, Y. (1976). *Computational Semantics*, E. Charniak and Y. Wilks (eds.). Oxford: North Holland.

Wilks, Y. (1978). Making Preferences More Active, *Artificial Intelligence* **11**.

Wilks, Y. (1988). Foreword to *Lexical Ambiguity Resolution: Perspectives from Psycholinguistics, Neuropsychology and Artificial Intelligence*, Small, S. I, G. W. Cottrell and M. K. Tanenhaus (eds.). San Mateo, CA: Morgan Kaufmann.

FIGURE CAPTIONS

Figure 1: The meaning constructors of the Conceptual Scaffolding.

Figure 2: Initial Scaffolding model of the metaphor "Chomsky rebuilt modern linguistics from the ground up", couched in a semantic case-frame representation. Key: Up and Down arrows represent orientation metaphors applied to a specific facet of a complex concept, whilst black-anode arrows represent actual causality and grey-anode arrows represent enablement.

Figure 3: The Conceptual Scaffolding of a sentence can be constructed compositionally, by a process of unification over local scaffolding structures stored in each lexical item. (Grey lines indicate the flow of information via unification; numbers in boxes indicate co-reference indices, or variables).

Figure 4: A dormant linkage between the concepts MUSICAL-INSTRUMENT and MUSKET is deemed to provide a plausible match hypothesis in the metaphor COMPOSER AS GENERAL when it becomes a domain cross-over for two competing waves of activation originating from the tenor (COMPOSER) and vehicle (GENERAL) matriarch nodes.

Figure 5: Cross-domain mappings produced by Sapper to reconcile the domains of Composer and General.

Figure 6: The Sapper Algorithm, as based on the exploitation of cross-domain bridge-points in long-term memory.

Figure 7: The Triangulation Rule (i) and the Squaring Rule (ii) augment the knowledge base with additional dormant bridges (shown as dashed lines), precompiled pathways that may later be used to form cross-domain analog bindings. The relational label “Attr” denotes a conceptual association/attribution, while the label “M” denotes an active conceptual bridge, i.e., a recognized metaphor.

Figure 8: Partial Sapper Memory Network linking the concepts Composer and General.

Figure 9: Scout — the retrieval component of Sapper— searches for a potential vehicle for the tenor concept Hacker. This figure illustrates that Scout has found seven different well-formed pathways to the concept node Surgeon, a central concept in the medicine domain, suggesting that this pairing might produce a highly systematic metaphoric mapping.

Figure 10: The Scout Algorithm

Figure 11: Two scaffolding structures built for the sentence “America withdrew its ambassador from the Russian puppet government of Bulgaria.”

Figure 12: Fragment of Long-Term Memory that contributes to the metaphors “Russia is a Puppet Master” and “Bulgaria is a Puppet”.

Figure 13: Cross-domain mappings produced by Sapper to reconcile the domains of Russia and Puppet-Master.

Figure 14: Fragment of Memory that contributes to the metaphors “Code: Lock” and “Password: Key”.

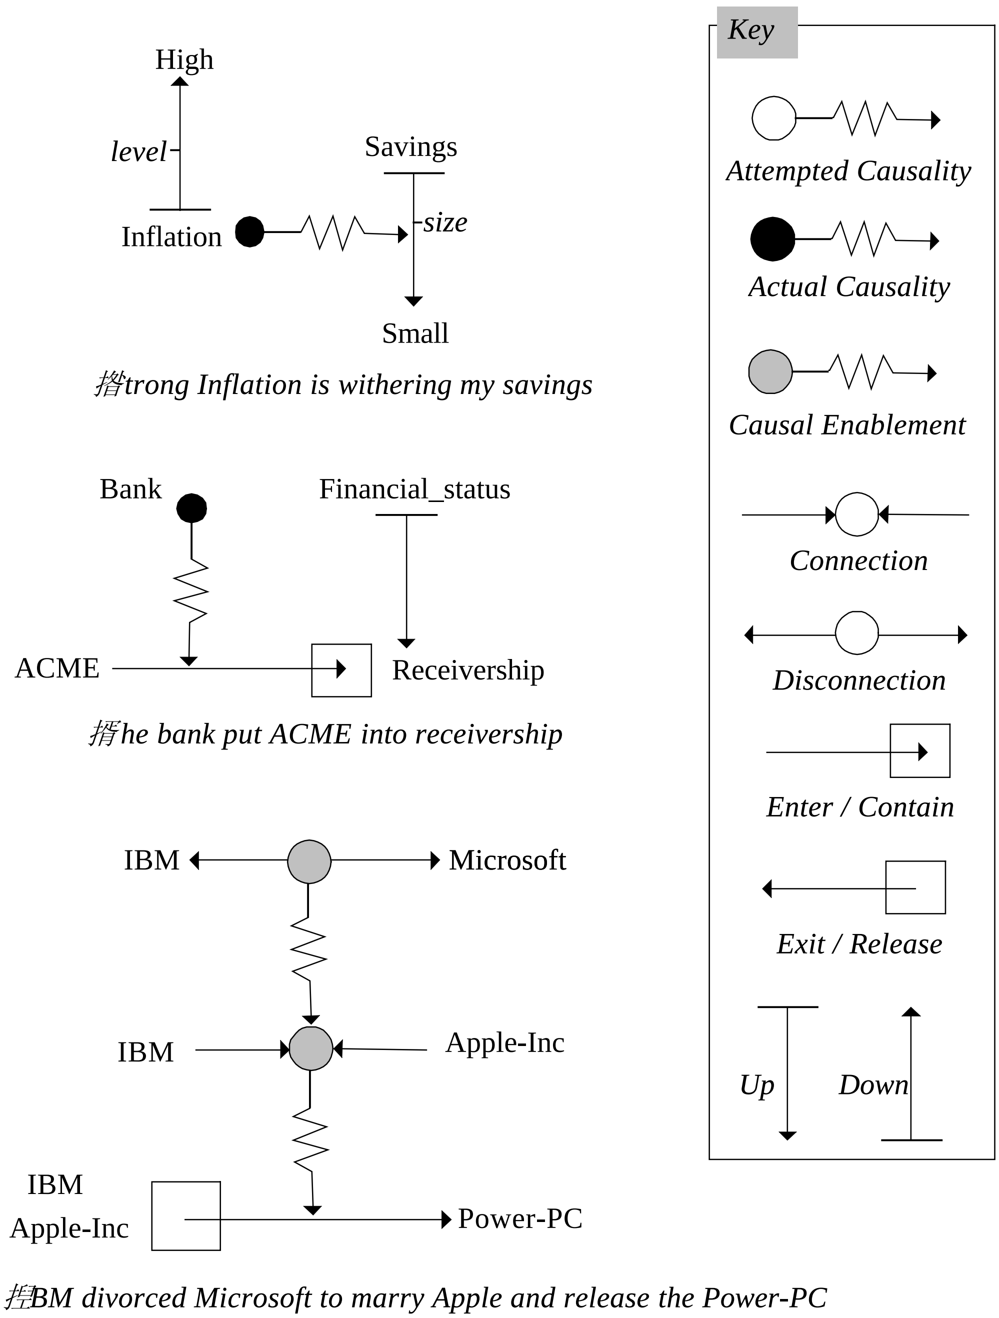

Figure 1: The meaning constructors of the Conceptual Scaffolding.

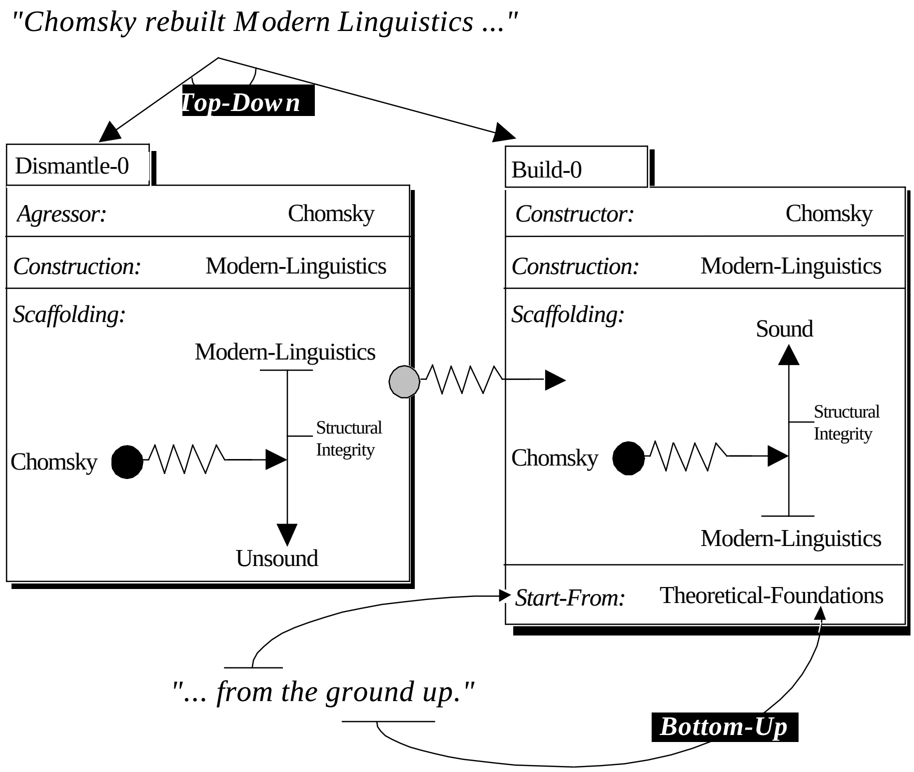

Figure 2: Initial Scaffolding model of the metaphor "Chomsky rebuilt modern linguistics from the ground up", couched in a semantic case-frame representation. Key: Up and Down arrows represent orientation metaphors applied to a specific facet of a complex concept, whilst black-anode arrows represent actual causality and grey-anode arrows represent enablement.

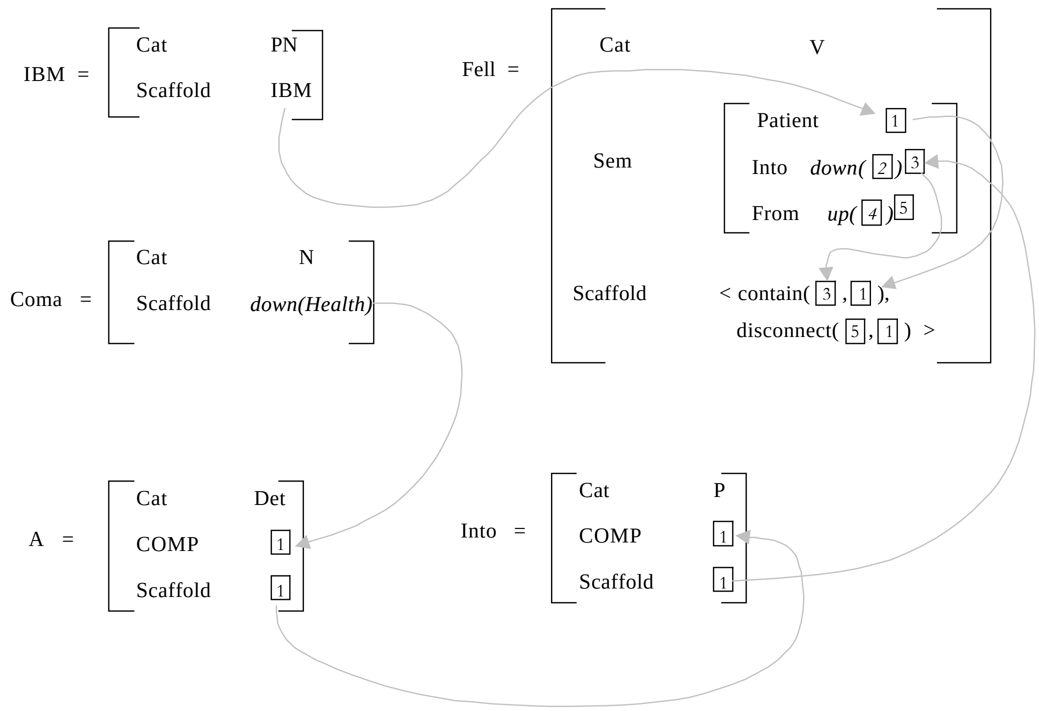

Figure 3: The Conceptual Scaffolding of a sentence can be constructed compositionally, by a process of unification over local scaffolding structures stored in each lexical item. (Grey lines indicate the flow of information via unification; numbers in boxes indicate co-reference indices, or variables).

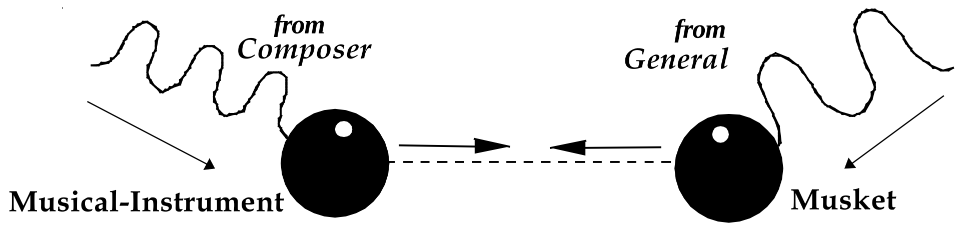

Figure 4: A dormant linkage between the concepts MUSICAL-INSTRUMENT *and* MUSKET *is deemed to provide a plausible match hypothesis in the metaphor* COMPOSER AS GENERAL *when it becomes a domain cross-over for two competing waves of activation originating from the tenor (*COMPOSER*) and vehicle (*GENERAL*) matriarch nodes.*

[.86]	If **Composer** is like **General**
[.25]	*Then* **Concert-Theatre** is like **Battle-Theatre**
[.75]	*and* **Orchestra** is like **Army**
[.94]	*and* **Musician** is like **Soldier**
[.98]	*and* **Musical-Instrument** is like **Musket**
[.95]	*and* **Baton** is like **Sabre**
[.92]	*and* **Musical-Score** is like **Battle-Plan**
[.93]	*and* **Percussion** is like **Artillery**
[.96]	*and* **Drum** is like **Cannon** *etc.*

Figure 5: Cross-domain mappings produced by Sapper to reconcile the domains of Composer and General.

When interpreting the apposition of two concepts T and V do the following:

Spread Activation from node T in long-term memory to a horizon H

Spread Activation from node V in long-term memory to a horizon H

When a wave of activation from T meets a wave from V at a bridge $T':V'$ *linking the tenor domain concept* T' *to the vehicle domain concept* V' *then:*

Determine a chain of relations R that links T' *to T and* V' *to V*

If R is found, then the bridge $T':V'$ *is balanced relative to T:V, so do:*

Generate a partial interpretation π *of the metaphor T:V as follows*

For every tenor concept `t` *between* T' *and T as linked by R do*

Put `t` *in alignment with the equivalent concept* `v` *between* V' *and V*

Store π *in a temporary work area for later use.*

Once all partial interpretations $\{\pi\}$ *have been found within horizon H, do*

Evaluate the richness of each interpretation π' *(e.g., count the number of mappings in* π'*, and add on individual similarity scores for each mapping)*

Sort all partial interpretations $\{\pi\}$ *in descending order of richness.*

Choose the first (richest) interpretation Γ *as a seed for overall interpretation.*

Work through all other partial interpretations π' *in descending order of richness*

If it is coherent to merge π' *with* Γ *(i.e., without violating 1-to-1ness) then*

$\Gamma \leftarrow \Gamma \cup \pi'$

Otherwise discard π'

When $\{\pi\}$ *is exhausted,* Γ *will contain the overall Sapper interpretation of T:V*

Figure 6: The Sapper Algorithm, as based on the exploitation of cross-domain bridge-points in long-term memory.

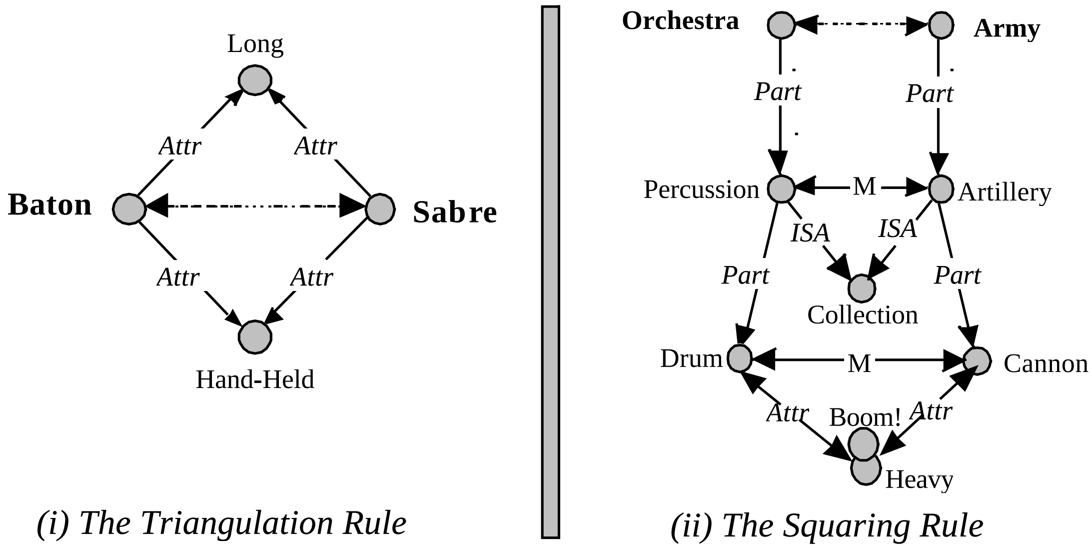

Figure 7: The Triangulation Rule (i) and the Squaring Rule (ii) augment the knowledge base with additional dormant bridges (shown as dashed lines), precompiled pathways that may later be used to form cross-domain analog bindings. The relational label "Attr" denotes a conceptual association/attribution, while the label "M" denotes an active conceptual bridge, i.e., a recognized metaphor.

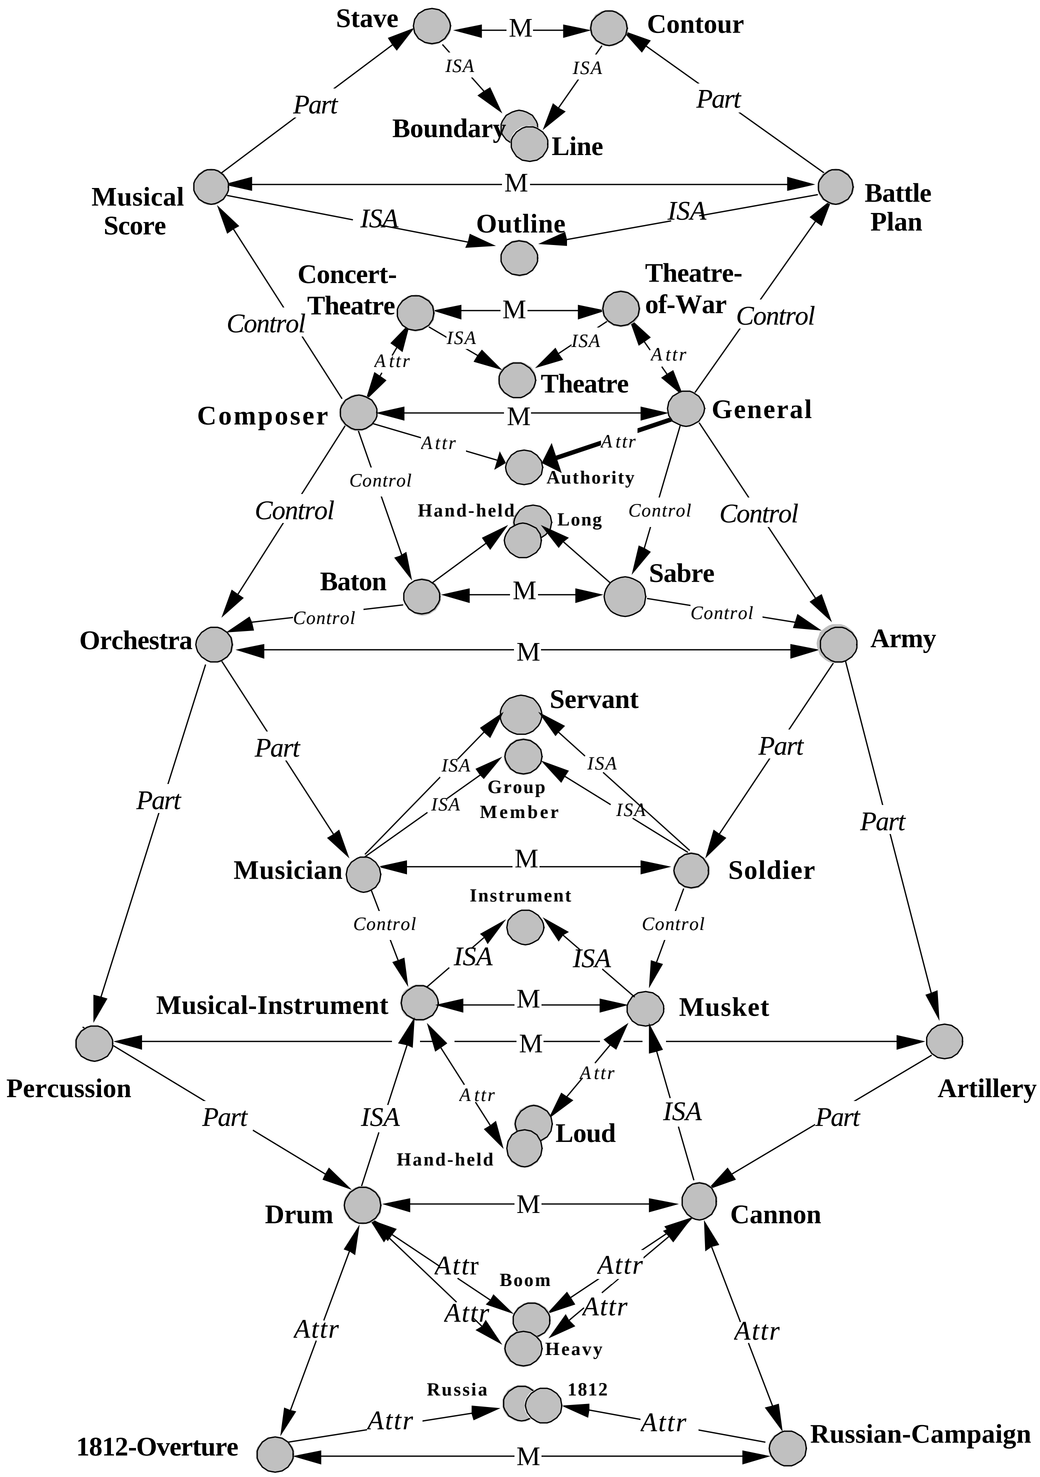

Figure 8: Partial Sapper Memory Network linking the concepts Composer and General.

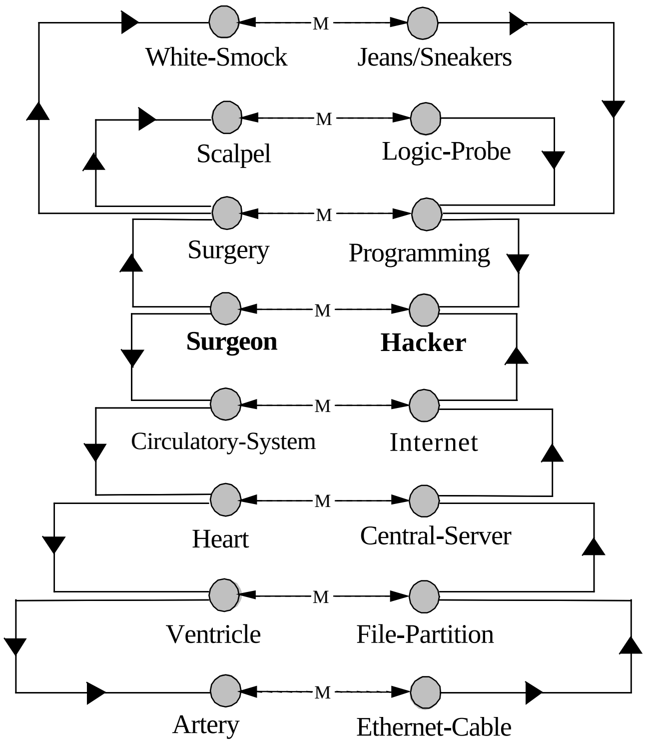

Figure 9: Scout — the retrieval component of Sapper— searches for a potential vehicle for the tenor concept Hacker. This figure illustrates that Scout has found seven different well-formed pathways to the concept node Surgeon, a central concept in the medicine domain, suggesting that this pairing might produce a highly systematic metaphoric mapping.

Spread Activation from node T (the Probe) in memory to a horizon H

When a wave of activation from T visits a node T′ at a bridge T′:V′ Then:

- Determine a chain of relations R (where $|R| \leq H$) linking T′ to T
- For all R, determine the node V that is reached if R is unfolded from V′
- If R can be unfolded to reach V, Then:
  - Generate a partial interpretation π of the metaphor T:V as follows:
    - For every tenor concept t between T′ and T as linked by R do
      - Align t with the equivalent concept v between V′ and V
      - Thus, $\pi \leftarrow \pi \cup \{<t:s>\}$
    - $\Phi \leftarrow \Phi + \pi$

Find $\Gamma \in \Phi$, the most mapping-rich pmap found within the horizon H

Return V, the counterpart of T in Γ, as the retrieved result.

Figure 10: The Scout Algorithm

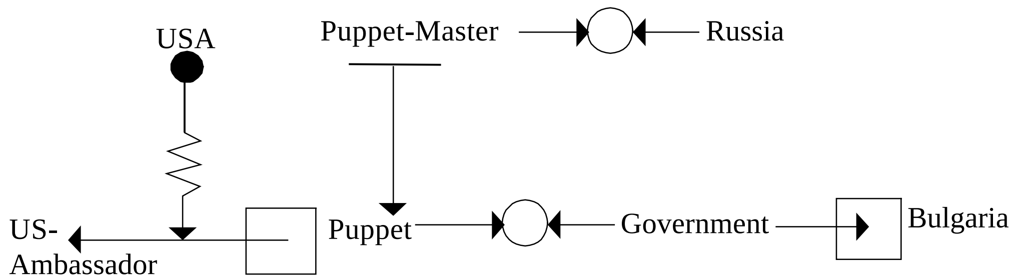

Interpretation (i): America withdrew its ambassador from the Bulgarian government which is a puppet of Russia

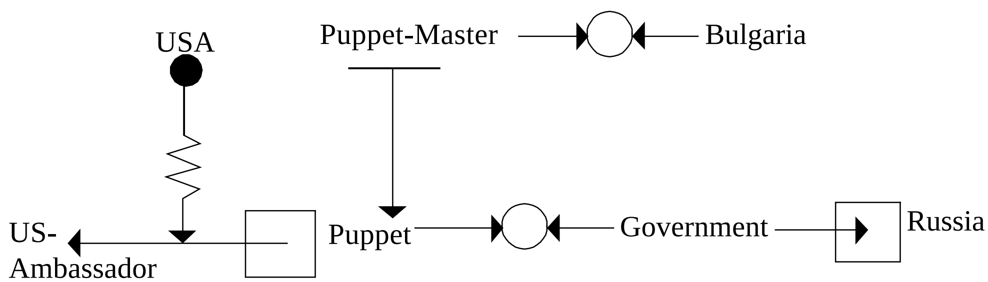

Interpretation (ii): America withdrew its ambassador from the Russian government which is a puppet of Bulgaria

Figure 11: Two scaffolding structures built for the sentence "America withdrew its ambassador from the Russian puppet government of Bulgaria."

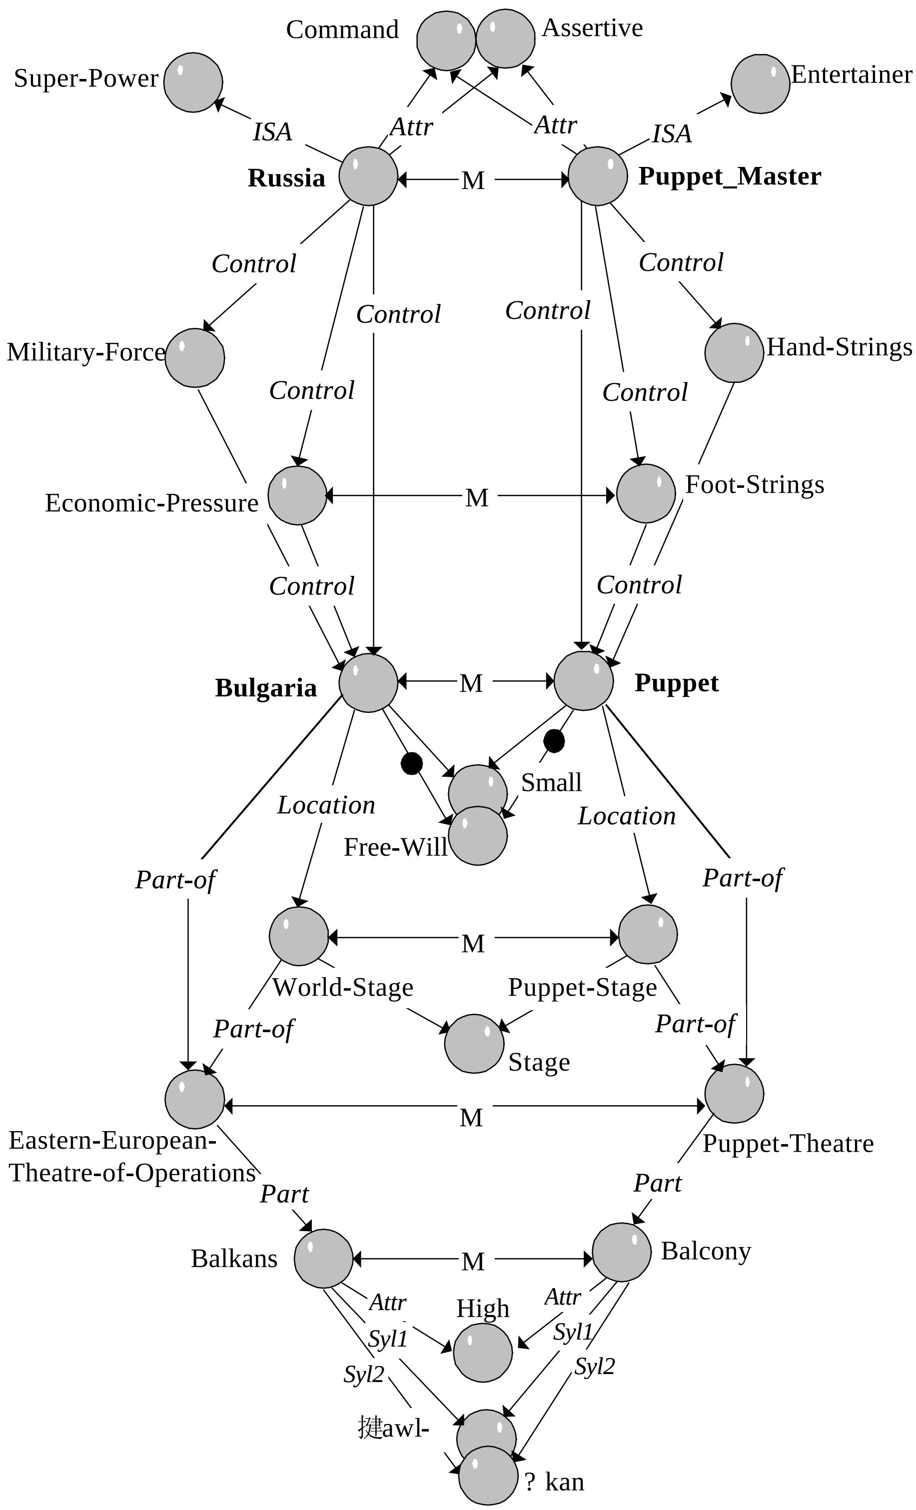

Figure 12: Fragment of Long-Term Memory that contributes to the metaphors "Russia is a Puppet Master" and "Bulgaria is a Puppet".

[.86] If **Russia** is like **Puppet-Master**
[.93] *Then* **Military-Force** is like **Hand-Strings**
[.75] *and* **Economic-Pressure** is like **Foot-Strings**
[.94] *and* **Bulgaria** is like **Puppet**
[.96] *and* **World-Stage** is like **Puppet-Stage**
[.95] *and* **East-Euro-Theatre-Ops** is like **Puppet-Theatre**
[.92] *and* **Balkans** is like **Balcony**

Figure 13: Cross-domain mappings produced by Sapper to reconcile the domains of Russia and Puppet-Master.

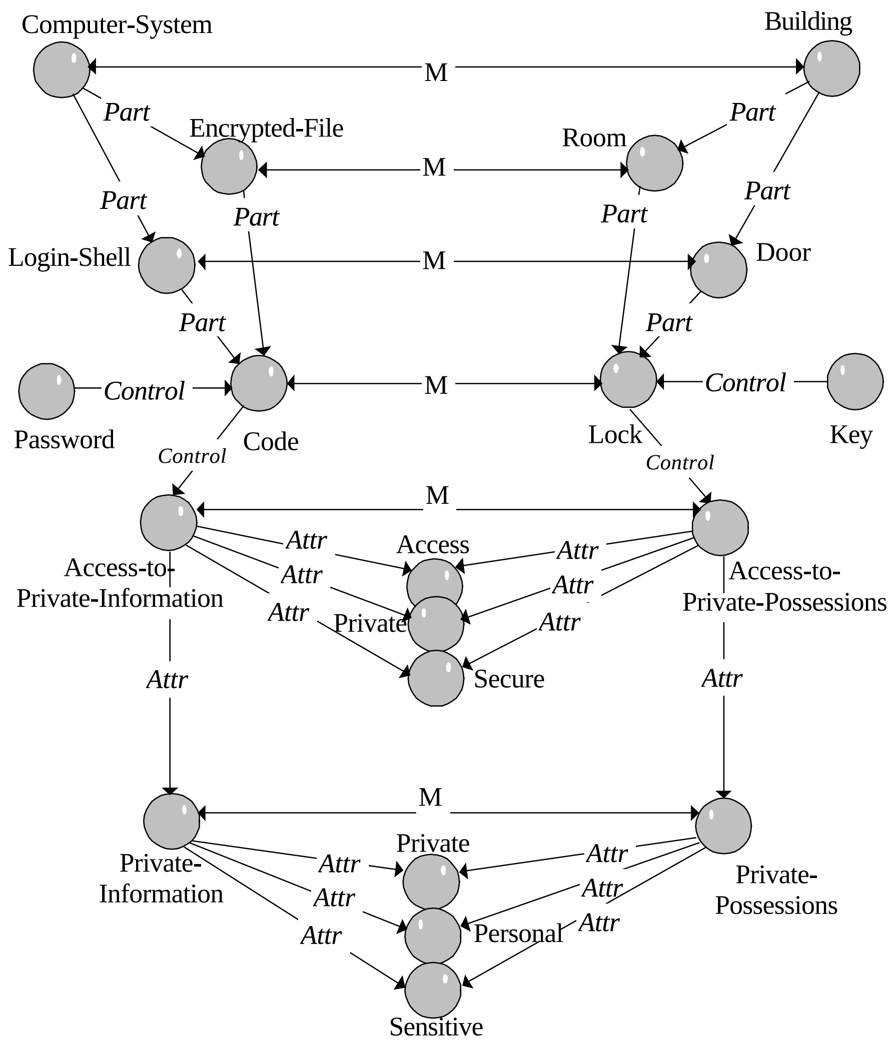

Figure 14: Fragment of Memory that contributes to the metaphors "Code: Lock" and "Password: Key".