Pragmatic Forces in Conceptual Integration:

How Blend Recruitment is Used to Make Metaphors More Compelling

Tony Veale,
School of Computer Applications,
Dublin City University,
Dublin, Ireland.
Email: tonyv@compapp.dcu.ie

Abstract

Metaphor and analogy are cognitive tools which, in serving specific communicative goals and descriptive needs, are subject to a host of pragmatic pressures. Knowing that these pressures will shape the interpretation of a given metaphor, an effective communicator will exploit them to structure the conceptual content of the metaphor in such a way as to maximise its perceived aptness and argumentative force to the recipient. This paper considers the form that such pressures can take, and the computational strategies that a communicator can employ to maximise the effectiveness of a given metaphor. This consideration will address not only computational models of structured metaphor/analogy and conceptual blending, but computational accounts of the role, and importance, of similarity in metaphor. We choose as our domain of discourse a collection of visual metaphors which highlights the effect of pragmatic strategies on metaphoric communication.

1. Introduction

A multitude of pragmatic pressures interact to shape the generation and interpretation of creative metaphors, analogies and blends. These pressures range from the need to relax strict isomorphism when identifying a mapping relationship between the tenor and vehicle domains, to recruiting intermediate blends, or self-contained metaphors, as mediators between certain cross-domain elements that would otherwise be considered to distant in conceptual or imaginistic space to make for an apt and aesthetically coherent metaphor. To apply the terminology of Hofstadter *et al.* (1995), such pressures fall under the broad rubric of *'conceptual slippage'*, since certain pragmatic accommodations must be made in the conceptual structure of a domain to facilitate a rich and meaningful interpretation. Slippage mechanisms allow individual ideas to fluidly shift from one underlying concept to another, in such a way as to enhance the structural coherence of an entire network of related ideas.

This paper considers the complex interaction between these various slippage pressures, and what they tell us about metaphor as a pragmatically motivated cognitive phenomenon. Though such fluid aspects of metaphor can be accounted for structurally, they nevertheless demonstrate that metaphor entails more than a simple structure-matching solution to the graph-isomorphism problem, harnassing a range of on-the-fly reasoning processes that can create complex transformational chains between entities. In most computational models of structural metaphor and analogy, such as the SME (Structure-Mapping Engine) of Falkenhainer, Forbus and Gentner (1989) and ACME (Analogical Constraint Mapping Engine) of Holyoak and Thagard (1989), two cross-domain entities are said to be analogical counterparts if they occupy the same relative position in their respective semantic structures. In contrast, the metaphors studied in this research suggest that analogical equivalence is much more than a matter of structural isomorphism: not only must two cross-domain concepts occupy the same relative semantic position, there must be some compelling semantic rationale for one to be mapped to the other.

Though such models have recognized the importance of pragmatic guidance in generating and interpreting metaphors, they frequently adopt a tangential approach to the problem (see Holyoak and Thagard 1989; Forbus and Oblinger 1990). Typically, it is assumed that such pragmatic guidance comprises a *fait accompli*, inasmuch as that certain potential mappings are earmarked as desirable or mandatory *before* the mapping process commences. The core of the computational model then attempts to construct a systematic lattice-work of cross-domain mappings that relates the tenor domain to that of the vehicle while including those earmarked preferences. The conceptual reasoning that actually yields this set of mapping preferences, or that generates the pragmatic force which pressurizes the algorithm into seeking a structural accommodation of both domains via some form of slippage, is not explicitly modelled. That this reasoning process has not been the subject of adequate computational inquiry is perhaps due to another issue in metaphor and analogy that has also been approached rather tangentially by the computational literature, and that is semantic (as opposed to wholly structural) similarity.

The phenomena studied in this paper help to make clear exactly what structural forms the conceptual rationale underlying the process of structural accommodation must take, and how this rationale is informed by considerations of semantic similarity. We demonstrate that in many cases the rationale is blend-centered, and that complex visual metaphors often recruit conventional visual blends to pragmatically motivate the key mappings of a larger metaphor. Because these blends represent established metaphors in their own right, they lend an immediacy to the metaphor in which they are incorporated, helping to make this encompassing metaphor eye-catchingly apt. But as we shall further show, we do not need to posit a new theory of metaphor to account for these phenomena, as the mechanics of this recruitment process are readily explained within the computational framework of Veale and Keane's (1994; 1995;1996;1997) Sapper model.

Given these goals, the paper assumes the following structure: section two considers, through the use of real-world examples culled from the political news media, some of the more interesting but computationally-vexing pragmatic pressures that collude to make a metaphor both informative and apt. Working within the general framework of conceptual blending theory (see Fauconnier and Turner 1994, 1998; Turner and Fauconnier 1995), this section suggests some schematic interpretations of these pressures that might in some sense be considered generic, and thus computationally attractive. Section three then sketches an overview of the operational principles of five current models of computational metaphor, focussing on one model, Sapper, in depth, and presents arguments as to why Sapper is best placed to accommodate the principles of conceptual blending. This accommodation requires us to introduce a new space into the theory of blending, dubbed *'constructor space'*, which in section four provides the means to computationally explain the pragmatic phenomena under discussion. The issue of semantic similarity is then addressed in section five; our discussion will be informed not only by these pragmatic phenemona, which patently require a computational model of similarity, but by the philosophical stance assumed by existing models of metaphor such as Sapper. The paper concludes with a summary and closing arguments in section six.

Before we proceed, it is important to clarify the terminological basis of our discussion. The phenomena under consideration in this paper all involve the conceptual integration of multiple mental spaces or conceptual domains, and rely on some mechanism of structure mapping to achieve this integration. Since the paper approaches these phenomena from the perspective of a computational model of metaphor that provides such a mechanism, we find it expositionally convenient refer to these phenomena under the umbrella term of *'metaphor'*, since it is our claim that the same computational mechanism is at work whether the phenomenon is best classed as a metaphor, analogy, literal concept combination or counterfactual blend. Originating as it does with Aristotle, the term 'metaphor' has the greatest claim to umbrella status (indeed, Aristotle himself cites proportional analogy as a sub-branch of metaphor; see Hutton, 1982).

2. The Problem: Pragmatic Issues in Metaphor Comprehension

The example of Figure 1 below represents a very real and complex illustration of the pragmatic pressures that interact to create a visually apt metaphor. Here we see the Economist newspaper use an easily identified piece of consumer gadgetry, a 'Tamagotchi' virtual pet, to make a searing indictment of the Japanese financial system: '*Firms such as Yamaichi* [Japan's 4th-largest brokerage, recently collapsed] *have been kept alive as artificially as the "virtual pets" in* Tamagotchi *toys: thank goodness those infernal gadgets are finally being turned off*'.

Figure 1: A striking visual blend of a 'Tamagotchi' game and the current Japanese financial situation after the recent Yamaichi Brokerage scandal. (Source: cover of the 'Economist', November 29, 1997)

Taken from a serious political newspaper — the 'Economist' — such a cartoon metaphor must be eye-catching yet appropriate, and complex (with a non-trivial political message) yet instantly understandable. In this section we discuss a variety of the cognitive phenomena responsible for the attention-catching potency of such complex metaphors, whilst also sketching some general schematic approaches to these phenomena in terms that can later be elaborated in the context of a specific computational model of metaphor.

2.1. Conceptual Blending

Much of what is commonly termed a metaphor in everyday language and the media can more accurately be described as a *blended* conceptual structure. In the *conceptual blending* or *many-space* theory of Gilles Fauconnier and Mark Turner, a blend is defined to be a selective integration of two or more conceptual structures to create another, a structure which owes its semantic foundations to its inputs but which also possesses an independent conceptual reality of its own (see Fauconnier and Turner 1994, 1998; Turner and Fauconnier 1995). The conceptual blending theory posits a multi-space extension of the classic two-space model of metaphor and analogy, in which the traditional inputs to the mapping process, the tenor and vehicle domains, are each assumed to occupy a distinct mental space, while the product of their conceptual integration is also assumed to occupy a separate output space of its own. This allows the newly blended concept to acquire associations and conventions that do not strictly follow from the logical makeup of its inputs. For instance, the concept *BlackHole* is a convenient and highly visual blend of the concepts *Blackness* and *Hole*, one that enjoys continued usage in scientific parlance despite evidence that blackholes are neither hole-like or black in any real sense.

In addition, the conceptual blending model allocates a distinct space, called *generic space*, to those schemas which underly, and frequently guide, the construction of a blend. These schemas operate at a low-level of description, typically the image-schematic level, and can serve both as selectional filters and basic structure combinators for the input spaces. From a constructive viewpoint, generic space can be seen as containing a set of basic conceptual schemas that enable the generation of new metaphoric and analogical mappings. Following Veale and Keane (1997), we denote the most fundamental of these schemas using the following path notation: *X—metaphor→Y*, which notes that the concept *X* can be connected via the schema *—metaphor→* to the concept *Y* when *X* and Y share a sufficient quantity of semantic structure. Conventional computational accounts of metaphor and analogy attempt to construct a self-consistent and systematic lattice of mappings by instantiating this schema for as many elements of the tenor and vehicle domain as possible; for instance, when interpreting the metaphor Surgeon as Butcher, it is consistent to instantiate this schema as *Scalpel—metaphor→Cleaver*, *Patient—metaphor→Carcass* and *Surgery—metaphor→Slaughter*.

We argue that for every distinct pragmatic force that can affect the shape of a given metaphoric mapping there may exist a corresponding mapping schema in constructor space. An understanding of constructor space and its likely conceptual contents may thus yield a computationally felicitous schematic account of the pragmatic forces that affect both a metaphor's generation and its interpretation. Equally, since blends are almost always conceptual constructs of convenience, either allowing greater imaginistic/visual access to an abstract idea, or extending the reach of existing inference procedures, a variety of important pragmatic forces may only be describable in terms of how they cause specific blends to be constructed or dissolved on the fly for the sake of communicative convenience. In the discussion that follows then, conceptual blending is a key element of our vocabulary for analysing the battery of pragmatic forces that can affect metaphor.

2.2. Recruitment of Sub-Metaphors and Blends

Complex metaphors have an internal structure which is itself frequently constructed from other, related metaphors. These sub-metaphors are ‘recruited’ (in the sense of Fauconnier and Turner 1994) to mediate between cross-domain elements of the larger metaphoric context. For instance, the clichéd *‘the pen is mightier than the sword’* can readily be recruited as part of a domain reconciliation between the concepts Author and General. In effect, the sub-metaphor *Pen—metaphor→Sword* acts as a foundation, or conceptual *bridge*, between the domains of Author and General around which a more elaborate interpretation can be constructed.

Following Veale and Keane (1993), we use the term 'bridge' to denote an instantiation of the schema X—*metaphor*→Y that supports further mappings at a higher-level (e.g., *Critics—metaphor→Enemies* and *Readers—metaphor→Army*). In this case, we describe *Pen—metaphor→Sword* as a perceptual bridge, since Pen can be seen as a metaphoric Sword for wholly perceptual reasons (both are long, pointed narrow and hand-held). Frequently however, pivotal elements of a metaphor will not directly share obvious perceptual qualities; for instance, the Yamaichi brokerage has no perceptual resemblance to a *Piggy-Bank*, but on a more abstract level, both can be seen as reservoirs of *Money* and aids to financial investment.

Sub-metaphors are recruited then whenever a lower-level metaphoric bridge is used as a foundation for higher-level comparisons. In perhaps the most interesting cases, however, the sub-metaphor is not itself apparent, as one's attention is instead focused on a blend structure that mediates between the elements of the sub-metaphor. This blend typically exhibits certain properties, both sensory and causal, in common with the elements it mediates between, even though those elements themselves may have little or nothing in common. The Tamagotchi *'piggy bank'* of Figure 1 is a clear example of such a mediating blend, inasmuch it connects two very disparate concepts (*Yamaichi* and *Electronic-Puppy*) via the concept *Piggy-Bank*. The schema that relates these concepts to each other is more complex than the basic X—*metaphor*→Y we have considered so far, in effect comprising a daisy-chain of these basic-level schemas in the form X—*metaphor*→Blend—*metaphor*→Y. But when the mediating blend is particularly apt, as is Piggy-Bank in this case, one does not explicitly notice the chain of schemas involved, but instead focuses on its pivotal element, the recruited blend.

2.3. Iconicity

A compelling metaphor will often exploit conventional iconicities of a domain, making the comparison both relevant and visually striking. This of course is particularly apt for visual metaphors, as seen in political cartoons, advertising imagery, and so forth.

But what makes a vehicle concept iconic? Iconicity is perhaps best described formally from a visual, information-processing perspective. For instance, one typically expects the pictorial expression of an iconic concept to exhibit a certain cleanness of line and minimality of detail, whereby a dearth of visual detail is used to evoke a wealth of conceptual associations — where iconicity is concerned, less is frequently more. For example, people readily associate cowboys with Marlboro cigarettes, Santa Claus with Coca-Cola, and Lions with MGM, provided some basic visual constraints are respected, namely the cowboys are rugged, Santa is fat, jolly and red, and the lions are fierce yet majestic. The Tamagotchi example of Figure 1 exemplifies this expectation of minimalism, since the potency of the *Piggy-Bank* image lies in stark contrast to the crude resolution with which it is etched. But since a discussion of visual aesthetics is outside the scope of this paper, we concentrate on a more conceptual and schematic aspect of iconicity here, one that is both intuitive and readily amenable to algorithmic exploitation.

For the purposes of the current work, we take iconicity to imply not only an established association from a tenor concept to a vehicle concept, but also a strong reciprocation of this association. For instance, though Ireland is commonly associated with rain, one does not necessarily think of Ireland when one thinks of rain. In contrast, Ireland is strongly associated with concepts such as *Guinness* (a black stout beer), *Shamrock* and *Leprechaun*, and these concepts in turn evoke strong associations of *Ireland*. In this respect, putting aside their visual properties, these concepts are iconic of the concept *Ireland*. Likewise, as shown in Figure 1, the hand-held game *Tamagotchi* is currently considered highly iconic of modern Japan, making this concept a good vehicle to describe the failure of the Yamaichi stock brokerage. In addition, Tamagotchi has strong associations with other concepts which are themselves strongly established icons of modern Japan, namely *Electronic-Gadget* and *Miniaturisation*. Our reciprocality view of iconicity is respected in these examples also, since all three concepts are mutually evocative of each other.

2.4. Double-Think in the Recruitment of Blends

It is an established premise of the computational literature that for a metaphor to work coherently, it must involve a systematic mapping of elements from a vehicle domain onto a corresponding set of elements in the tenor domain (e.g, see Black 1962; Winston, 1980). For instance, in mapping an iconic concept such as *Guinness* onto the social class structure of modern Ireland, one might map the creamy white *Head* onto the concept *Political-Class* (i.e., ‘the fat-cats’), and the black *Beer* below this onto the concept *Working-Class* (i.e., ‘the great unwashed masses’). However, such mappings may frequently need to be mediated by an intermediate concept, for either structural, aesthetic or pragmatic reasons. On this point, much of the computational literature is silent.

For example, when one describes a person as a *Wolf*, one rarely employs a realistic schema for *Wolf*, but a stereotypical model that many people now know to be false. This archetype is closer in nature to the cartoon caricatures of Chuck Jones (e.g., lascivious, treacherous, ruthless and greedy) than to accepted reality (e.g., that a wolf is a family animal, with strong social ties). This caricature is an anthropomorphic and highly visual blend of properties drawn from both *Person* and *Wolf*, which allows a cognitive agent to easily ascribe human qualities to a non-human entity. More importantly, the recruitment of such blends facilitates a fundamental conceptual role of metaphor which, pace George Orwell’s ‘1984’, we term *‘Doublethink’*, namely, the ability to hold two complementary perspectives on the same concept in mind at the same time, and to combine or blend these perspectives for reasons of inference when necessary.

Consider again the Tamagotchi visual metaphor of Figure 1. The creators of this metaphor (the Economist magazine, November 29th 1997) exploit the Japanese associations of the Tamagotchi game to describe the situation facing Japan's banking regulators after the downfall of the Yamaichi stock brokerage. The metaphor particularly stresses the options open to the regulators – to prop up (i.e., 'feed') the ailing brokerage, or let it fail (i.e., 'die'), while viewing the whole financial fiasco as a 'game' gone wrong. Tamagotchi games conventionally focus upon electronic pets such as puppies or kittens, which the player (the regulator?) is supposed to nourish and nurture via constant interaction. This central animal is thus a good metaphor for Yamaichi, but the visual impact would clearly be diminished if the artist simply substituted a picture of a bank, no matter how iconic, into the game. This is thus a situation in which direct mapping between tenor and vehicle elements lacks a sufficient pragmatic force of its own.

Fortunately, a blend is available, that of a *Piggy-Bank*, that possesses the necessary iconic features to substitute for both Yamaichi and the Tamagotchi puppy in the metaphor. A *Piggy-Bank*'s strong associations with money and savings make it an ideal metaphor for Yamaichi, while its visual appearance makes it an obvious (after-the-fact) counterpart to the electronic animal of the game.

This is where the notion of 'double-think' applies. While being a metaphor for both a brokerage and a puppy, the *Piggy-Bank* blend is allowed to exploit contradictory properties of both. Most obvious is the orientation of the piggy-bank – its 'belly-up' position is an iconic visual commonly associated with animals – indicating that Yamaichi is either already bankrupt (dead) or seriously insolvent (dying). This inverse orientation would make no sense if applied to a literal image of a bank, yet it is perfectly apt when applied to another artefact, the piggy-bank, due its blend of animal visual properties (the most important here being 'legs'). Following what Medin and Ortony (1989) term the thesis of *'psychological essentialism'*, is seems that people freely transfer the animistic essence of *Pig* to *Piggy-Bank* on the grounds that common perceptual features are often indicative of deeper causal properties. The *Piggy-Bank* concept is not simply a structural substitute then for Yamaichi and puppy, but a 'living' blend of both.

2.5. Recasting

In the case of the Tamagotchi metaphor of Figure 1, the pragmatic situation is actually even more complex than this. Though the concept *Piggy-Bank* is identified as an appropriately visual mid-point between a financial institution and a puppy for the metaphor to work, recall that the source of this key sub-metaphor is not actually a puppy at all, but an electronic simulation of one. We thus introduce the idea of a resemblance schema, taking the form *X—resemble→Y*. A resemblance relation is simply an elaboration of the schema *X—metaphor→Y* for relating concepts whose correspondence is determined by the sharing of perceptual (i.e., appearance-related) properties. By daisy-chaining these schemas together, we can generate a host of more complex *transformational schemas* for relating concepts that considerably removed in imaginistic space. The transformational chain linking Yamaichi to the Tamagotchi puppy is thus: *Yamaichi—metaphor→PiggyBank—resemble→**Pig**←metaphor—Puppy←resemble— TamagotchiPuppy.* In effect, the electronic Tamagotchi puppy is as much (or as little) a real dog as a porcelain piggy-bank is a real pig. Both domain concepts simply need to be recast for the metaphor as a whole to cohere.

Figure 2: A bowling metaphor is used to convey the rough-and-tumble of modern Russian politics.

Indeed, recasting seems to be a structural phenomenon which is key to stamping visual coherence on a metaphor. Consider for instance another graphic metaphor from the cover of the 'Economist' (November 22, 1997), which illustrates the rough-and-tumble dynamism of modern Russian politics. To convey the main thrust of the magazine's leader column, namely that certain once-prestigious Russian politicians continue to suffer humiliating downfalls while Boris Yeltsin remains upright and stable throughout, the 'Economist' chooses a bowling metaphor in which different pins represent various politicians, and in which bowling balls represent the fickleness of public opinion.

The metaphor, illustrated in Figure 2, is well-chosen not only because bowling is a populous sport associated with the general public as a whole, but because the up/down/stable/rocking status of the pins conforms to a conventional mode of discourse in politics. However, visual coherence cannot be bought simply by painting the faces of the politicians involved onto the appropriate pins, as the conceptual and imaginistic distance between bowling pins and people is such that the result would simply look contrived. Instead, the cover's creator uses not bowling pins but nested Russian dolls, of the political variety one finds in tourist shops (a common doll features a *Yeltsin* encasing a *Gorbachev* encasing a *Breshnev* and so on).

While possessing an iconic visual quality, such dolls also resemble both bowling pins *and* politicians, and so act as a perfect mediating blend between the end-points of the metaphor.

2.6. Internal Recruitment

A blend which is recruited to act as a mapping intermediary in this way also acts a visual *precedent*, in effect grounding the mapping in shared background knowledge between creator and reader as well as securing the aptness of the mapping. However, not all elements of the metaphor may be externally grounded in this fashion. For instance, in the case of the Yeltsin bowling cartoon, the Russian finance minister Anatoly Chubais is also illustrated using a Russian doll/bowling pin blend, yet there is no background precedent for this. Nevertheless, there exists an *internal* precedent – Boris Yeltsin. Because Yeltsin is also depicted in this fashion, and because Chubais is a strong analogical counterpart of Yeltsin (both are powerful male Russian politicians), it makes sense that any grounding applied to Yeltsin can also be analogically transferred to Chubais. So while Yeltsin visually maps to the first bowling pin via the transformational chain *Yeltsin—resemble→YeltsinRussianDoll←resemble—BowlingPin#1*, Chubais maps to the second via the chain *Chubais—metaphor→Yeltsin—resemble→YeltsinRussianDoll ←resemble—BowlingPin#1←resemble—BowlingPin#2*. It seems from such examples that metaphor can possess an incestuous quality, feeding not only off other metaphors and blends recruited from outside, but also upon itself.

2.7. Analogical Inferencing

Analogy can be seen as a didactic form of metaphor in which the purpose of communication is to educate by comparison. However, while many metaphors are simply descriptive, with aesthetic rather than educational goals, metaphors can also possess a *take-home* message which the reader transfers from the vehicle domain to the tenor. For instance, in comparing Japan to a Tamagotchi, the Economist's take-home message is the opinion that perhaps the Japanese government has viewed the problems of financial regulation as a game, while treating favoured institutions like Yamaichi as *'virtual pets'*. This form of transfer-based inferencing is easily incorporated into models of analogy and metaphor such as SME, ACME, LISA and Sapper, given that the cross-domain mapping established by these models acts as a *substitution-key* which dictates how elements of the vehicle domain can be rewritten into the tenor domain.

However, not all metaphors provide a sufficient key for transferring elements of the vehicle into the tenor. For instance, in the Russian bowling metaphor, what is to made of the fact that certain political *kingpins* are shown falling on their sides? This idea of a *'fall from grace'* has a strong metaphoric history in politics, conventionally denoting failure due to scandal, but this is a metaphor that must be recruited *from outside* the current context rather than identified and exploited internally. So, when presented with an image of a falling Chubais doll/pin, one must draw upon political knowledge associated with a 'fallen' analogical counterpart of Chubais from outside the current context, if it is not already appreciated that this particular politician is in a perilous position. For instance, one might defer to Nixon, another politician, and his political fall, via the analogical chain <u>*Chubais*</u>*—metaphor→Nixon* ***—perform→Resignation*** *—metaphor→*<u>*Fall*</u>. In essence, we simply need to find a path that metaphorically links the concept Chubais to the concept Fall, and this path should contain the semantic sub-structure to be analogically carried into the tenor domain; in this case the connecting sub-structure suggests that Chubais might ***perform*** an act of ***resignation***. It is necessary that the cognitive agent reason via an analogical counterpart like Nixon, rather than via some fixed association of the concept Fall, since the concept Fall may have different metaphoric meanings in different contexts (e.g., one would not infer that a falling share-price should also resign).

2.8. Determining a Relevant Scope

As already noted, iconicity plays a key role in determining relevance in visual metaphors. Because Japan is currently associated in a strong sense with the Tamagotchi game (many might use the word 'guilt' in this context), while Tamagotchi itself is clearly evocative of Japan (as many miniaturised electronic gadgets are), Tamagotchi therefore acts as a deeply iconic metaphor for modern Japan.

But the metaphor of Figure 1 is clearly driven by the need to communicate the current economic situation in Japan as it particularly applies to the Yamaichi scandal. Indeed, knowing little about Yamaichi itself, many readers would be hard-pressed to recognise any iconic associations with what was until recently a rather anonymous Japanese brokerage. There thus exists a strong pragmatic pressure to widen the scope of the metaphor, in this case to Japan as a whole, while insisting that any metaphor chosen to reflect Japan will encompass Yamaichi in a recruited sub-metaphor.

This enlargement of the metaphoric context serves two pragmatic goals: firstly, the larger ‘Japan’ metaphor serves to place Yamaichi and its woes in a given cultural context, while secondly, it structurally enriches the metaphor by allowing more cross-domain elements to participate (in this case, Tamagotchi itself, which is iconic of Japan but not Yamaichi).

The encompassing context can be chosen in much the same way as iconic vehicles are chosen for a metaphor. For instance, Yamaichi is strongly associated with Japan, while Japan is itself causally related to Yamaichi via its regulators. In contrast, though Yamaichi is associated with mountains (its name means ‘Mountain Number 1’), the concept Mountain is not saliently associated with Yamaichi. Thus the concept Japan, rather than a concept like Mountain, can be recognised as providing a larger metaphoric context in which to work. Once Japan is chosen to act as the new tenor of the metaphor, the tentative vehicle Tamagotchi can then be chosen for its iconic value. It remains for the creative agent to ‘run’ the metaphor of ‘*Japan is a Tamagotchi game’* with the caveat that Yamaichi receives a cross-domain mapping in the interpretation. Computational models of analogy and metaphor such as SME, ACME and Sapper each provide for this pragmatic directive.

Such a situation is evident in another cartoon cover from the Economist, this time its May 1997 China Review, in which China’s reticence to open itself to global markets is explored. The particular focus of the cover is thus the rather delicate issue of insularity and xenophobia as it is perceived in China’s dealings with Western powers. To emphasise these issues, the Economist cleverly enlarges the metaphoric context to one in which China is fore-grounded as an actor in a theatrical setting, essentially the *world stage*, allowing the topic of xenophobia to be humorously communicated as a form of *stage-fright*. The concept World-Stage is thus the enlarged tenor of the metaphor, but this in turn allows China to be recruited as a nervous actor with first-night anxiety. A visual blend, that of a ceremonial Chinese-Dragon, is also recruited to act as an iconic intermediary between the concepts China and Actor (see Figure 3).

Figure 3: Visual Blends employed to give animate characteristics to a Nation. China is shown both as a Panda-Bear and as a theatrical Chinese-Dragon, as the context dictates.

The result is particularly apt, since the concepts China and Chinese-Dragon share the associations Pride, Aggression, Ceremony, Chinese-Culture and Chinese-History, while Chinese-Dragon and Actor share many common associations of the theatre, such as Performance, Make-Believe, Make-Up, Special-Effect and Stage-Prop. Compellingly, the concepts Xenophobia and Stage-Fright also share a range of associations such as Anxiety, Fear and Stranger. The final image is a humorously effective one then, that of a Chinese dragon shyly poking its head through the curtains of a theatrical stage.

3. Computational Models of Metaphor and Conceptual Blending

The previous section has sketched a variety of related pragmatic forces in metaphor that cannot simply be dismissed as peripheral phenomena, as each force we have examined clearly contributes in a strong sense to the overall meaning of the metaphors and blends in question. However, most computational models of metaphor, even those that have ostensibly been designed with a view to modelling pragmatic pressures, have little to say on many of these issues, while no model currently provides a satisfactory account of all of these pressures working together.

In this section then, we outline the basis of five current models of computational metaphor, and highlight what scope each possesses to fully address these issues.

3.1. Structural Isomorphism: A Computational Basis for Metaphor and Analogy

At the heart of analogy and metaphor lies a structure-mapping process that is responsible for creating an isomorphic correspondence between semantic sub-structures of the tenor and vehicle domains. Isomorphism is a mathematical notion that guarantees the systematicity and coherence of any resulting interpretation, by ensuring that each relation and object of the tenor domain receives at most one correspondence in the vehicle domain. Isomorphism is central to metaphor and analogy because, in logical/computational terms, all meaning is expressed via structure; if a cognitive process does not respect structure, it cannot respect meaning, and thus, cannot itself be a meaningful process. Though a graphtheoretic mathematical notion, isomorphism is implicit in the writings of many non-mathematical philosophers of metaphor; Black (1962), for example, describes metaphor as a process in which a blackened sheet of glass inscribed with translucent markings (the vehicle) is placed over a visual scene like the night sky (the tenor). Only those stars which show through the markings are thus visible to the observer, and in this way a sub-graph isomorphism between glass and scene is created. For instance, a darkened glass inscribed with a picture of a winged horse might only allow light to shine through from those stars that comprise the Pegasus constellation.

3.2. SME: The Structure-Matching engine

As described in Falkenhainer, Forbus and Gentner (1989), SME—the *Structure Mapping Engine*—occupies one extreme of a functional continuum, and may be described as an exhaustively *optimal* and *maximal* approach to structure mapping. SME tirelessly produces all possible interpretations of a given analogical pairing, each alternate interpretation deemed maximal in the sense that no additional correspondence can be added to it without destroying its internal *systematicity* and coherence. Additionally, SME is optimal in the sense that it scores each alternate interpretation, and indicates the best mapping according to a predefined systematicity metric.

Actually, SME is a *configurable analogy toolkit*, capable of applying different match rules to different mapping tasks. Heuristic modifications to SME are additionally reported in Oblinger and Forbus (1990), which replace the factorial merge stage of the original algorithm with a new sub-optimal greedy merge. However, as argued in Veale *et al.* (1996a), even with such modifications SME is fundamentally unsuited to the mapping of structures in which richly detailed character/object descriptions—as opposed to high level causal actions—play an important role.

In SME parlance, a systematic collection of inter-structure correspondences is termed a *gmap* (global mapping). Initially, a set of *kernel* or *root gmaps* is constructed by systematically comparing the corresponding arguments of identical predicates in each structure. This set is grist for the core of SME, a combinatorial process which then produces successively larger combinations of these partial gmaps (called *pmaps*) until maximal global mappings are generated. Clearly, the size of the initial root set is a key factor in the tractability of the combination process; SME employs the notion of *structural support* to limit the size of this set, exploiting systematicity across the nested organization of predications in each structure as an evidential basis for generating new roots. However, we demonstrate in Veale *et al.* (1996a,1996b) that this support is not at all visible to SME in object/character-based metaphors, which tend to involve many shallow, tree-structured representations linked via common leaves, rather than a few, deeply nested tree representations linked via common superclasses.

3.3. ACME: The Analog Constraint Matching Engine

ACME, the *Analogical Constraint Mapping Engine*, also places great emphasis on the property of mapping *systematicity*, or *isomorphism*, but eschews the exhaustively optimal and maximal strategy pursued by SME. Instead, ACME constructs a constraint network for each new analogical problem to model the various pressures of similarity, context and isomorphism which shape the final interpretation. This network is the subject of a parallelized constraint relaxation process, from which a sole interpretation emerges, one that is neither guaranteed to be optimal, or maximal, or, for that matter, even wholly systematic. Unlike SME, ACME guarantees *nothing*, embodying a heuristic rather than complete approach to the problem. Indeed, ACME pursues what may be called a *natural* or *evolutionary* model of computation, in which environmental forces pressurize a system into converging toward a *good*, rather than optimal, solution (much like the *CopyCat* model of Hofstadter and Mitchell 1988; Hofstadter *et al.* 1995).

Like SME, ACME is a structure matcher which compares two domain descriptions in a predicate-calculus-style representation. Hierarchical structure in such descriptions— which is originally expressed via nesting of predications—is translated into a series of inhibitory and excitatory linkages in the ACME network. Nodes in this network correspond to possible entity correspondences between the source and target domains; once the network is activated, the activation levels of these nodes gradually converge toward asymptotic values as the network proceeds through a succession of epochs before eventually settling. An ACME network is deemed to have settled when a certain large proportion of its nodes have reached their asymptote. Yet while neither maximal or optimal, ACME is slower than SME, and is certainly less systematic; this contrary result is borne out in the empirical analysis of Veale *et al.* (1996a,b).

3.4. LISA: Learning, Induction and Schema Abstraction

The LISA model of Hummel and Holyoak (1996) can be seen as a conceptual descendant of Holyoak and Thagard's ACME, inasmuch as it too is a connectionist model of structure-mapping, albeit one that casts a thrifty eye on the scale of its network representations. Whereas ACME constructs a large, specialized network whose number of neurons/nodes is the square of the number of concept symbols in each domain, and whose number of inter-node connections is a fourth power of this domain size, LISA employs the notion of synchrony of firing to reduce the scale of its network considerably. Rather than dedicate a network node to representing a role binding between a specific pairing of a predicate symbol and an argument symbol, LISA instead assumes that the nodes representing these symbols are related if they are both firing (emitting an output signal) *in phase*.

Another feature which characterizes the evolution of LISA from ACME is its postulation of a semantic layer of micro-feature units (representing generic concepts such as Male, Action, etc.) in which higher-order structural representations are grounded. It is this aspect of LISA that makes it most interesting from the perspective of pragmatic-dictated similarity as discussed in section 2. LISA thus uses feature-based semantic criteria to judge if two conceptual structures are similar, in addition to the isomorphism-based criterion of structural similarity employed in SME and ACME. In effect, this layer of semantic features serves as a simple form of generic space against which to perform a metaphoric mapping, by providing the common semantic vocabulary with which to synthesise the input domains.

3.5. Tabletop: A Micro-domain for the Exploration of Similarity

Mining the same micro-domain vein as Hofstadter and Mitchell's (1988) Copycat architecture, the Tabletop model of Hofstadter and French (1995) and French (1995) is a non-deterministic study of the role of similarity and analogy in high-level perception. Tabletop is an environment for studying the pragmatically shifting nature of similarity judgements that one experiences during the analysis of *'do as I do analogies'*, in which a *student* attempts to replicate the actions of a *master* given certain, contextually-enforced restrictions. For instance, imagine a colleague with a broken left arm immovably encased in plaster of paris, who attempts to point to the exact bone in his elbow that is damaged; unable to point directly to his left elbow (covered as it is in plaster), he uses his left hand to point to the equivalent spot on his right elbow, exclaiming 'here's the bone I broke'. People generally have no trouble in comprehending this kind of analogical action; in fact, the mapping may be so transparent as to be unremarkable to the hearer.

Figure 4: 'Henry' points to a coffee-cup in the Tabletop micro-domain. 'Eliza' must now choose the corresponding piece of tableware to touch in response.

Tabletop employs this kind of 'perceptual analogy under contextual constraints' to the micro-domain of a coffee-shop table-top, on which various restaurant paraphernalia have been arranged. On one side of the table sits the master (denoted 'Henry' by Hofstadter and French), who touches a single item on the table; on the opposite side sits the student (denoted 'Eliza', to complete the Pygmalion metaphor), who attempts to mimic this action from her own perspective. Such a situation is illustrated in Figure 4.

Eliza chooses what she considers the best analog to Henry's choice from a host of perceptual gestalts, each informed and shaped by different pragmatic 'grouping pressures'. For instance, if Henry touches his coffee cup, as illustrated in Figure 4, Eliza may respond by touching her own coffee-cup. However, this rather literal reaction ignores both the geometric position of Henry's choice (her right versus his middle) and the high-level grouping of objects in which Henry's coffee cup participates (knife to the right, glass to the far right, coffee-pot to the left). Eliza may thus find that her plate of chocolate cake better reflects the geometry and chunking of Henry's cup, and touch that accordingly.

Hofstadter and French model these various slippage pressures in a non-deterministic, stochastic manner, by embodying each pressure via a computational device termed a *'codelet'*. Individual codelets represent potential choices and actions, and are given as much attention as their numeric ranking of urgency dictates. Tabletop thus moves its attention amongst a variety of different codelets, each competing at differing levels of urgency for the system's attention. Because Tabletop's attention mechanism is probabilistic over these urgency levels, different runs of the system may produce different results, yet on average, the most plausible results are produced most often.

Though restricted to a specific micro-domain, Hofstadter and French's system is nonetheless a model of considerable clarity that serves to pin-point those forces that shape an analogical/metaphoric interpretation. In an important sense, the model we describe in the next section, Sapper, is very much informed by the Tabletop model.

3.6. Sapper: A Bridge-Building Model of Structure-Mapping

Like SME and ACME, Sapper is a computational model of metaphor and analogy founded upon the notion of structure-mapping between domains (see Winston 1980; Gentner 1983). However, unlike SME and ACME, but like Tabletop and LISA, Sapper requires that two cross-domain concepts have more in common that an isomorphic structural setting if they are to be paired in an interpretation of a given metaphor. In addition to structural isomorphism, Sapper requires that two analogical concepts either share a common set of features (abstract or concrete) or be structurally related to another pair of analogical concepts that do. Concepts that share a number of semantic features or attributes are said to be linked by a *'bridge relation'*, and it is upon such *'bridges'* that Sapper grounds the interpretation of a given metaphor. For instance, the concepts *Scalpel* and *Cleaver* share the associations *Sharp*, *Blade* and *Blood*, and thus a bridge relation is established between both. Higher-level analogical correspondences can be founded upon this bridge if the corresponding concepts relate to the bridge in an identical semantic fashion; thus, because *Surgeon*s use *Scalpel*s, and *Butcher*s use *Cleaver*s, a mapping between *Surgeon* and *Butcher* can be grounded in the bridge relation between Scalpels and Cleavers. Bridges based upon low-level literal and perceptual similarities, such as *Sharp*ness, correspond to basic attributive metaphors, and are considered by Sapper as instantiations of the basic mapping schema *X—metaphor→Y*. Sapper views metaphor interpretation as a process of bridge-building in which new bridges are constructed using existing bridges as foundations; thus Sapper might construct the bridge *Surgeon—metaphor→Butcher* by building upon the lower-level bridges *Scalpel—metaphor→Cleaver* or *Surgery—metaphor→Slaughter*.

At the algorithmic core of Sapper lies a graph-matching process (see Veale and Keane 1997 for a full complexity analysis), one which exploits the bridge schema *X—metaphor→Y* to ensure that certain, pivotal elements of a cross-domain mapping are grounded in perceptual similarity. Sapper is also then, in an important sense, a partial theory of memory organization, inasmuch as it suggests that long-term memory is not simply a passive reservoir of information, but a reactive system that assimilates new conceptual structure by extending bridge relations to neighbouring structures that are, at a local level, semantically similar. Sapper thus employs a pro-active view of long-term memory in which shared associations between concepts are automatically recognised and noted, making low-level bridge construction a memory-centred rather than mapping-centred task. Built upon this reactive memory is a structure-mapping process that exploits these low-level bridges as construction cues for the elaboration of a global mapping between the tenor and vehicle spaces; this mapping, essentially a graph-isomorphism, serves as a semantic interpretation of a metaphor or analogy.

3.6.1. Cross-Domain Bridging in Sapper

In Sapper terminology, bridging schemas lay down bridge relations in memory between two concepts that are recognised to share some local regularity of structure. These bridges are initially dormant, since each represents a potential, rather than an actual, analogical correspondence between concepts; it remains for a structure-mapper to later confirm that a given bridge does indeed contribute to a systematic cross-domain interpretation. At this point a bridge becomes active. This *dormant / active* distinction proves very useful in lexical priming and spreading-activation applications of semantic memory, since an *active bridge* can carry activation (or pass markers) into metaphorically related domains (see Quillian 1968; Loftus and Collins 1974). Analogy and metaphor in Sapper thus have a representational effect on memory, actually warping the contours of long-term conceptual representation, rather than merely relying on memory as a passive warehouse of static domain descriptions.

The Sapper model typically employs two distinct constructors to augment its long-term memory representation with new dormant bridge relations—the *Triangulation Rule* and the *Squaring Rule*. The

Triangulation rule is invoked whenever two concepts share a common association or superclass; for instance, in a metaphor that relates surgeons to butchers, triangulation may occur among *HumanFlesh*: *Meat* and *Flesh*, *Scalpel*: *Cleaver* and *Sharp, WhiteSmock* : *Apron* and *Clothing*, and *OperatingTheatre*: *Abattoir* and *Location*, laying down dormant linkages between the schemata *HumanFlesh* and *Meat*, *Scalpel* and *Cleaver*, *whiteSmock* and *Apron*, and *OperatingTheatre* and *Abattoir*. In essence the triangulation rule is a formalization of a similar principle which underlies the plan recognition model of Hendler (1989), in which two high-level concepts can be seen as plan analogues 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 at airport customs. As noted previously, 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 low-level semantic features.

The *Squaring* rule is a *second-order* constructor that acts upon the linkages laid down between low-level feature concepts by the triangulation rule to build bridges between higher-level concepts associated with those similar but domain-incongruent features. For instance, it may be used to build (or *reinforce)* the bridges *Surgery : Slaughter*, *Scalpel: Cleaver*, and *Patient : Cleaver* in the *"surgeons are butchers"* metaphor. Sapper thus 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. The use of these constructors is illustrated in Figure 5.

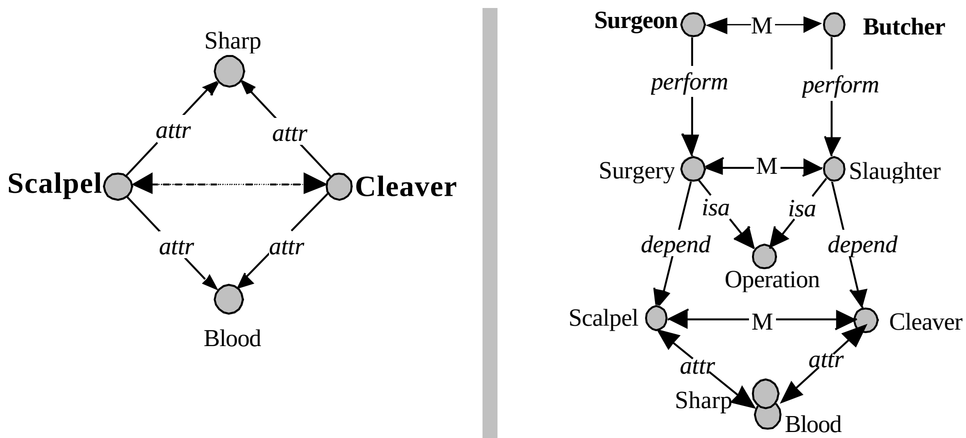

(i) The Triangulation Rule *(ii) The Squaring Rule*

Figure 5: The Triangulation Rule (i) and the Squaring Rule (ii) augment memory with additional bridges (depicted as $\longleftarrow m \longrightarrow$, a shorthand for the bridging schema), indicating potential future mappings.

We now turn to a consideration of how these local regularities of structure, expressed in memory as dormant conceptual bridges, are exploited by a structure-mapping algorithm.

3.6.2. Structure-Mapping in Sapper

The Sapper structure-mapping algorithm comprises two consecutive structure-building stages. The first of these performs a bi-directional breadth-first search from the root nodes of the source space or *vehicle* (denoted S) and target space or *tenor* (denoted T) in memory, to seek out all relevant bridges that may potentially connect both domains.

This search is constrained to occur within a maximum search horizon H (typically $H = 6$), to avoid the combinatorial explosion that occurs due to the considerable arboricity of complex concepts. This stage, a pseudo-code description of which is presented in Figure 6, produces an initial set of intermediate-level partial mappings (or *pmaps*) by aligning any isomorphic semantic pathways that meet at a cross-domain bridge.

Function Sapper::Stage-I (T:S, H)

Let $\Pi \leftarrow \varnothing$

Spread Activation from roots T and S in long-term memory to a horizon H

When a wave of activation from T meets a wave from S at a bridge T′:S′
linking a target domain concept T′ to a source concept S′ then:
Determine a chain of relations R that links T′ to T and S′ to S
If R is found, then the bridge T′:S′ is balanced relative to T:S, so do:
Generate a partial interpretation π of the metaphor T:S as follows:
For every tenor concept t between T′ and T as linked by R do
Align t with the equivalent concept s between S′ and S
Let $\pi \leftarrow \pi \cup \{t{:}s\}$
Let $\Pi \leftarrow \Pi \cup \{\pi\}$
Return Π, a set of intermediate-level pmaps for the metaphor T:S

Figure 6: The construction of intermediate-level pmaps is performed using a bi-directional search from the root nodes of both conceptual domains.

Sapper uses the same criterion of predicational matching as SME, namely: inter-concept relations (predicates) must match identically, while their arguments may not. Two semantic pathways are thus isomorphic if (i) they are of the same length; (ii) they comprise the same semantic relations, in the same order; and (iii) there exists a bijective mapping between the concept arguments of both. For instance, the source-domain pathway *Surgeon—perform→Surgery—depend→Scalpel—attr→Sharp* is isomorphic with the target path *Butcher—perform→Slaughter—depend→Cleaver—attr→Sharp*, and both combine to produce the pmap *{<Surgeon : Butcher>, <Surgery : Slaughter>, <Scalpel, Cleaver>}*

The second stage proceeds in much the same fashion as the greedy extensions made to SME in Forbus and Oblinger (1990): the most elaborate intermediate-level pmap is chosen as a seed mapping to anchor the overall interpretation, while other pmaps are folded into this seed, if it is consistent to do so, in descending order of the richness of those pmaps. Pseudo-code for this stage is presented in Figure 7.

Function Sapper::Stage-II (T:S, Π)

Once all partial interpretations $\Pi = \{\pi_i\}$ have been gathered, do:

Evaluate the quality (e.g., mapping richness) of each interpretation π_i

Sort all partial interpretations $\{\pi_i\}$ in descending order of quality.

Choose the first interpretation Γ as a seed for overall interpretation.

Work through every other pmap π_i in descending order of quality:

If it is coherent to merge π_i with Γ (i.e., respecting 1-to-1ness) then:

Let $\Gamma \leftarrow \Gamma \cup \pi_i$

Otherwise discard π_i

*When $\{\pi\}$ is exhausted, **Return** Γ, the Sapper interpretation of T:S*

Figure 7: The construction of global-level pmaps is performed using a seeding algorithm than coalesces intermediate-level pmaps in the order of their mapping richness.

3.6.3. Modelling Structural Slippage in Sapper

As is evident even from a micro-domain as simple as Tabletop's, context typically imposes a variety of interlocking pressures that necessarily complicate our decision-making processes. For instance, hould Eliza touch her own coffee-cup in response to Henry touching his, even though her cup occupies a distinctly different geometric position and gestalt, or should she *loosen* her conception of coffee-cup to encompass her plate of chocolate cake, which bears some superficial similarities to a coffee-cup but which, more importantly, occupies a similar relative position to Henry's designated item? Moving outside the closed world of the coffee-shop, and its limited conceptual repertoire (which Hofstadter and French term its *'Platobet'*), we see that any cognitive process that employs structure-mapping is equally likely to demonstrate this form of slippage. The ability to fluidly reorganize one's conceptual structures in response to contextual obstacles and pressures is thus a necessary element of metaphoric, analogical, and blend-centered mapping.

Interestingly, Sapper's squaring rule does, in an important sense, already serve as a recursive mechanism for conceptual slippage, one that allows two apparently dissimilar concepts to be reconciled if two other concepts, each related in the same way to the first pair, can themselves be reconciled. For example, *Surgery* can be seen as a form of *Slaughter* by virtue of being able to see a *Scalpel* as a type of *Cleaver*. But as defined, Sapper is inflexible in the face of another manifestation of conceptual slippage, a form which effects the structural interrelation of these concepts. For instance, two conceptual spaces or domains may be organized according to two similar, but superficially different, sets of conceptual relations, e.g., one space may make extensive use of the substance relation (e.g., *Table—substance→Wood*) while another uses the contain relation (*Body—part→Blood*), yet it might be desirable to treat these relations as having the same meaning for analogical purposes. Transitivity across such relations is also an issue of conceptual slippage: if we know that in one conceptual *domain A—cause→B—enable→C*, and that in another *X—cause→Z*, it might be pragmatically sensible to map A to X and C to Z. Drawing inspiration from the Tabletop model, we outline in this section a simple extension to the Sapper model that allows for these forms of slippage.

Underlying this *soft* extension of Sapper is a Tabletop-like *slipnet* in which different semantic relations (i.e., graph labels) are probabilistically connected (e.g., $P_{slip}(part \rightarrow contain) = 0.9$). Operating in conjunction with this slipnet is a semantically motivated structure-warping rule, which essentially softens the standard Sapper triangulation and squaring rules to allow non-isomorphic structures to be mapped. Sapper can thus comprehend analogies between domains that have been defined at different levels of detail and redundancy: for instance, in the *SportsCar* domain one might state that the *Engine* contains *Piston*s which *control* the *Wheel*s, or alternately, that the *Piston*s *control* the *Crankshaft* which in turn *control*s the *Wheel*s. When mapping this source structure then to that of either the *Jaguar* or *Puma* panther say (an analogy used by Ford for two of their sports cars), it may be necessary to either contract or stretch the target structure to accommodate the possible occurrence of the node *Crankshaft* (which might or might not map to *LegMuscle*, say).

Given two pmaps of equal depth (i.e., each composed of paths of a given length), a probabilistic *rigidity* measure of how much slippage each involves can be ascertained, as a product of the necessary slippage probabilities entailed by each. Thus, a pmap that maps X—$part$→Y—$contains$→Z to A—$part$→B—$contains$→C has a rigidity measure of *1.0*, while one that maps the same path to A—$contains$→B—$part$→C has a rigidity measure of $0.9 \times 0.6 = 0.54$. These measures can in turn be incorporated into a quality metric that prefers rigid pmaps over their looser variants that have *slipped.*

If $S \text{—} \ldots \rightarrow S_1 \text{—} R_1 \rightarrow S_2 \text{—} R_2 \rightarrow S_3$ *is a path being followed in the source space,*

and $P_{slip}(R_1, R_2) > \varepsilon$ *(a minimal rigidity threshold)*

Then

$S \text{—} \ldots \rightarrow S_1 \text{—} R \rightarrow S_3$ *is also a path that should be pursued in the source*

Where

$R = R_1$ *if* R_1 *is a causal relation, otherwise* $R = R_2$

Figure 8: The Core Slippage Principle employed in Sapper.

Given the existence of a relational slipnet to handle label slippage, the complementary problem of structural warping can be handled with the single, compositional rule of Figure 8. The action of this rule is simple yet effective: two successive semantic relations R_1 and R_2, linking two concepts S_1 and S_3 via an intermediary S_2, can be *snipped* to produce a path that links S_1 and S_3 directly; if R_1 is a causal relation (such as *cause, enable, support*, etc.) then it is favoured as the relation that directly connects S_1 and S_3; otherwise R_2 is chosen. If applied at every stage of a given pathway's development, this rule is capable of removing a significant number of linkages, as many are as needed to make the pathway structurally isomorphic with a mirror pathway in the target domain. For instance, *part→cause* reduces to *cause*, as does *cause→part,* while *part→substance* and *part→contains* both reduce to *part*. As illustrated in Figure 9, the concepts *Engine* and *CrankShaft* are temporarily removed from the source picture to accommodate a mapping between *Muscle* and *Piston*.

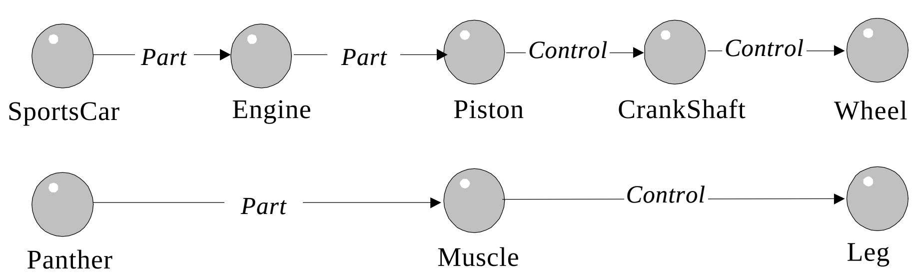

Figure 8: Path simplification in the SportsCar domain yields a path isomorphism with the Panther domain.

Note that this convenient deletion of *Engine* is indeed both temporary and non-destructive, inasmuch as it effects just this single pathway. Other pathways that ultimately find a mirror partner in the target domain may instead provide a mapping for $Engine$ (for instance, a pathway between $FuelCap$ and $ExhaustPipe$ will necessarily pass through $Engine$, mapping it to either $Brain$ or $Heart$). Ultimately, Sapper will choose those pathways, simplified or otherwise, that collectively contribute to the richest overall mapping. The issue then of whether slippage is warranted at any level in a mapping is effectively resolved by Sapper's in-built optimality pressure to construct the most coherent and elaborate interpretation possible.

3.7. A Computational Model of Conceptual Blending

As noted in section 2, a cognitive model which is to adequately capture key pragmatic dimensions of metaphor must account not only for the traditionally accepted characteristics of metaphor, such as cross-domain mapping, but also for the mechanics of blend recruitment that make this mapping possible. In this section then we consider which of the above computational models is most capable of accommodating the conceptual demands of Fauconnier and Turner's blending theory. The evidence, as we see it, points most clearly to Sapper as the architecture most directly informed by the theory of conceptual blending, specifically in respect of such issues as the organising role of generic space, the interaction of this space with the input spaces, the emergent character of blend space, and the conceptual integration that occurs between each of these spaces.

Consider first the notion of generic space, and the algorithmic guise this notion might assume in a computational model. Recall that the generic space of a blend contains those conceptual schemas that underlie, and serve to unite, the individual constructs of the input (e.g., tenor and vehicle) spaces, thus providing the common conceptual vocabulary of the blend. But as such, a model like SME makes little or no appeal to the idea of a generic space, primarily relying as it does on structural alignment to reconcile its inputs Though SME's mapping may be influenced by the surface properties of the concepts in each domain, it does not depend on these properties to generate a mapping, and thus does not seek to recruit intermediate blends that make the mapping more compellingly vivid. Likewise, neither does the ACME architecture support an obvious counterpart to generic space, as again, this model does not look beyond its input constructs to generate a mapping. In contrast, the Tabletop architecture makes recourse to a *slipnet*, or probabilistic network, of platonic concepts (collectively dubbed its *'platobet'*) to reason about those elements of a problem that are contextually similar to one another. Those *platobetic* concepts that become activated in the course of a mapping problem, such as that of Figure 4, can thus be construed forming as the generic space of the mapping. In a similar fashion, LISA makes recourse to a substrate layer of semantic nodes to ground the mapping of two higher-order structures, such that any node that becomes active corresponds to a semantic feature that is generic to both input spaces. Likewise, Sapper also grounds the mapping of two conceptual structures in literal similarity, via its use of the triangulation and squaring rules, so that those concepts which serve as the basis of triangulation (such as *Blood* and *Sharp* in Figure 5(i)) can, like their counterparts in LISA, be seen as forming the generic conceptual vocabulary of the mapping.

Though each of the five models supports an obvious counterpart to the input spaces of a blend, few actually give computational support to the idea of an independent blend space. Again, SME and ACME support no direct equivalent of a blend space, for since neither model supports an explicit model of semantic memory, both are powerless to describe the effects on semantic memory of a mapping, the most notable effect being the creation of a new conceptual space in which the blend is to reside. Similarly, the modification of semantic memory and the creation of new conceptual spaces (which might correspond to a structured partition of its slipnet) are beyond the remit of the Tabletop architecture. In contrast however, the LISA model makes specific, theoretical claims about the process of schema induction in metaphor and analogy, demonstrating how one can acquire, via abstraction, generic conceptual schemas from specific mapping problems. However, as the label 'generic' here implies, these abstractions will most likely correspond to the rarified contents of generic space rather than to the elaborate structures of blend space, since a primary function of blend space is to facilitate the accretion of emergent features around a newly blended structure.

We instead argue that Sapper best captures the notion of blend space via its use of *active conceptual bridges*. While dormant bridges serve to capture in memory the potential for combination between two input spaces, the newly awakened active bridges of a metaphor serve to explicitly represent the corresponding elements of these spaces that actually fuse to create the blend. For example, given the core bridge of a blend, such as that connecting *Composer* and *General*, it is a simple matter to employ the squaring rule in reverse to visit all those bridges with which that bridge is structurally consonant, such as *Army : Orchestra* and *Artillery : Percussion*.

Collectively then, the bridges constructed by Sapper for a given mapping (both dormant and active) correspond to the conceptual integration network that ties each of the contributing spaces together in a blend. Dormant bridges serve to relate elements of the input spaces to each other while simultaneously highlighting those elements of generic space that make the correspondence possible, whereas active bridges relate the input spaces directly to the resulting blend space. Since these bridges are represented in long-term memory as explicit traces of a specific blended space, access to a given bridge will, in turn, provide access to the concepts that connect to the end-points of this bridge. For instance, given the bridge *Cannon : Drum*, a cognitive agent can, by a process of *spreading activation* (see Quillian, 1968; Collins and Loftus, 1983; Charniak, 1983), retrace the efferent associations of its conceptual end-points *Drum* and *Cannon* back into the domains of *Composer* and *General*. This process, which Fauconnier and Turner dub *'unpacking a blend'*, is of particular cognitive and computational importance, since it underlies an agent's powers of blend introspection, decomposition and re-organization.

3.8. The Computational Role of Constructor Space

The conceptual foundations of blend theory are conventionally defined in terms of four spaces, namely a generic space, two input spaces and a blend space, though this basic architecture is easily extended to accommodate blends that fuse more than two inputs. However, to offer a detailed computational picture of the algorithmic processes employed in blending, we argue that one must posit a fifth, ontologically different type of space, one which we dub *'constructor space'*. While the four conventional spaces contain conceptual structures of varying degrees of abstraction and experiential grounding, this fifth space contains the computational rules of structure composition that allow these conceptual spaces to be structurally aligned and coherently fused.

These rules, or *constructors*, are inherent to all computational models of analogical and metaphoric mapping, though some models exploit constructors more transparently than others. In SME, for instance, one can employ different sets of explicit alignment rules for solving different types of mapping problem, while in contrast, the workings of ACME's constructors are altogether more implicit, manifest only in the way different mapping nodes of the constraint network are hard-wired to each other. In Sapper, these constructors correspond to the triangulation and squaring rules, and as such, they can be seen as responsible for the dynamic construction of the generic space of a blend.

The computational interaction between spaces under this algorithmic view is illustrated in Figure 9, where those inter-space dependencies that are computational in nature are depicted as grey arrows, while those that are wholly conceptual are depicted in black. Following this diagram, one can see how the constructor space uses specific rules of structural organisation to determine the conceptual middle-ground of a blend, and thus populate the generic space with the concepts necessary to achieve an integrated network.

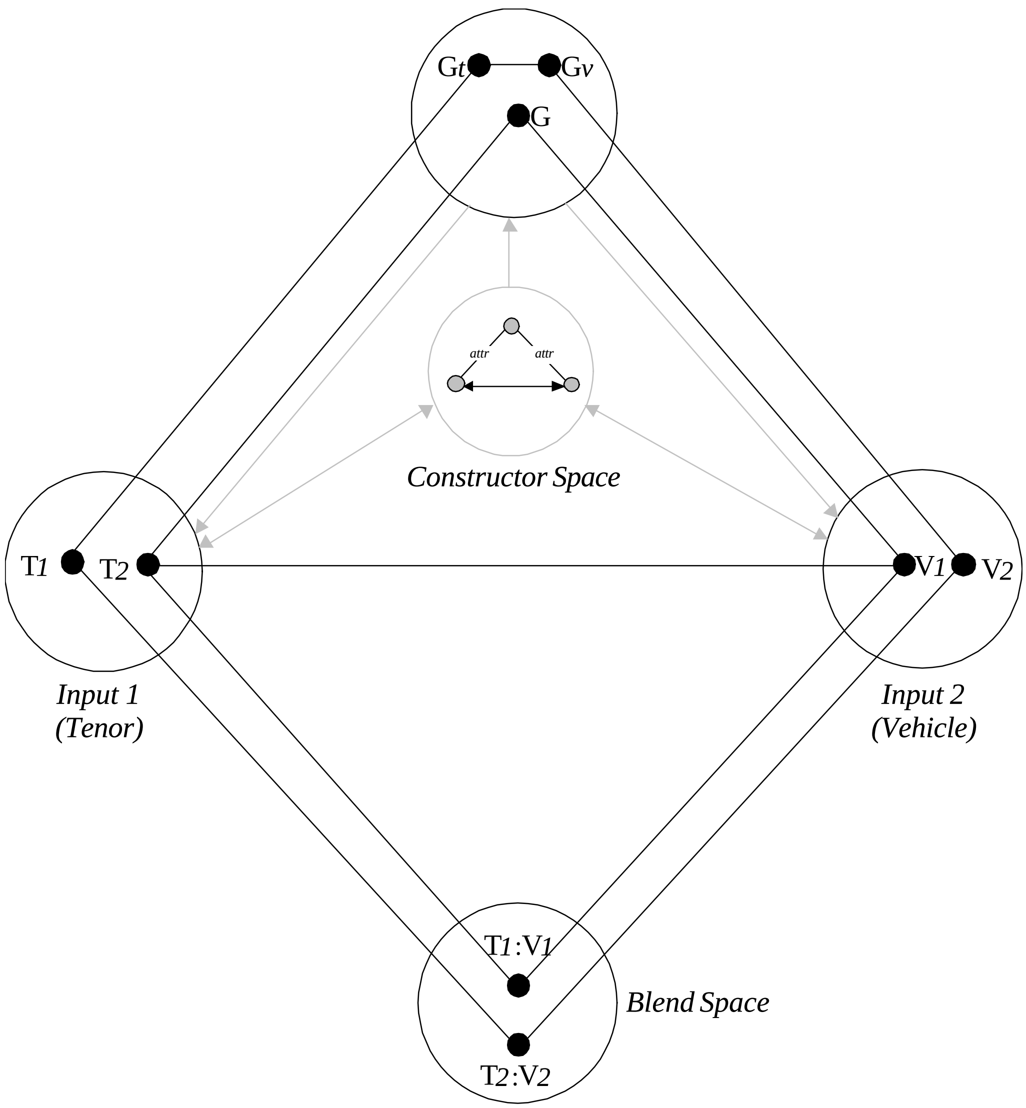

Figure 9: The Conventional 4-Space model of Blending, Augmented with a Fifth 'Constructor' Space when

In this particular example, the triangulation rule is used to reconcile the concepts T_1 and V_1 via a shared association *G*, causing *G* to become newly active in the generic space of the blend. Similarly, the squaring rule is used to reconcile the concepts T_1 and V_2 via an established, underlying metaphor $G_t : G_v$, causing this metaphoric schema to likewise become active in generic space. This latter mechanism can be seen at work in blends like *"Norman Mailer is Hemmingway crossed with Patton"*, where established metaphors such as *Pen as Sword* may be recruited to reconcile and blend the underlying schemas of *Author* and *General*.

The rules of constructor space simultaneously determine the generic content of the blend and apply this content to structurally reconcile the input spaces via a coherent isomorphic mapping. Therefore, only a selection of the possible mappings between the input spaces is chosen, since many of these mappings are unsystematic when considered in combination with others. For instance, when mapping *Surgeon* to *General* (the basis of the CNN/military metaphor *"surgical airstrike"*), one can map either *Enemy-Soldier* to *Cancer-Cell* or *Enemy-Army* to *Bacteria*, but not do both, since the latter does not cohere with the former. What is projected into the blend space then is a maximal collection of mutually systematic bridges, each bridge representing a fusion of counterpart elements from the input spaces. In Figure 9, input elements T_1 and V_1 are fused in the blend space as $T_1{:}V_1$, while T_2 and V_2 become fused as $T_2{:}V_2$. These bridges, newly activated and established in long-term memory, may later serve as the generic basis for an even more complex future blend.

We feel it necessary to posit the existence of an explicit constructor space, at the risk of complicating the already elegant four-space model, because of the apparent freedom one has in determining which constructors are most suited to a given blend. Indeed, the choice of constructors can sometimes be as important as the choice of the input spaces themselves. For instance, though the default tenants of Sapper's constructor space are the triangulation and squaring rules, an agent may feel it necessary to add an additional slippage rule, such as that of Figure 8, when the context demands that the principle of mapping isomorphism be relaxed somewhat. Indeed, some contexts actually call for the structure-preserving principle of isomorphism to be abandoned altogether, forcing an astute agent to populate its constructor space with rules that perversely destroy structure. Consider the mapping problem of Figure 10, in which an agent is asked to assign the letters A through J to the nodes of a graph such that no two adjacent nodes contain letters that are themselves alphabetically adjacent.

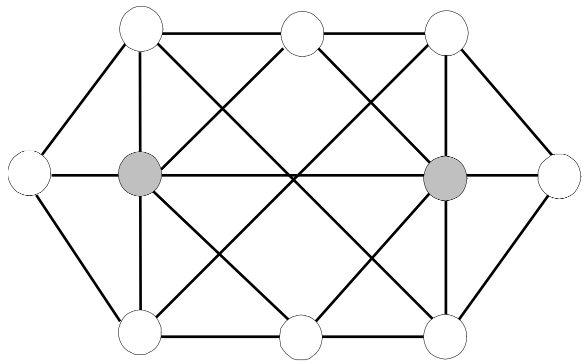

Figure 10: Map the Letters A through Z onto the Ten Nodes of this Graph Such That No Two Nodes Contain Alphabetic Neighbours.

Though the problem of Figure 10 is one of reasonable complexity (there are *10!* potential mappings), it is easily solved if one views it as a structure-destroying blend or formal *dis-analogy*. The alphabetic sequence A to J can itself be organised as a graph, albeit a very linear one, where A connects to B, B to C, and so on to J. The problem can thus be rephrased as follows: map this alphabetic graph onto the graph of Figure 10, such that the adjacency structure of the former is completely destroyed. With this as the goal, a key insight is to map A and J onto the grey nodes of Figure 10, since A and J are the least connected nodes of the alphabetic sequence, while these grey nodes are the most highly connected nodes of the problem graph. With A and J in place, the problem becomes so constrained as to be near trivial. This insight, if formalised as a structure-destroying constructor (or more accurately, a *destructor*) rule, can be applied in similar, future problems that demand such a perverse mapping.

In the following sections we shall increase the potential population of constructor space even further, by arguing that the pragmatic phenomena discussed in section 2 are best modelled via computational, rather than purely conceptual, schemata. Though algorithmic-level descriptions are typically not to be preferred over conceptual descriptions, save for purposes of demonstrating cognitive tractability, these phenomena are clearly best understood at this level, since they exist by virtue of an agent's efforts to make its metaphors and blends readily understandable to their audience. Thus, what constitutes 'understandable' in this context is predominantly a computational issue, as the likelihood of comprehension is closely related to the ease with which the intended interpretation can be computed at its target.

4. Constructor Space: Structural Schemas for Pragmatic Mapping

Viewed from a Sapper perspective, the basic schema *X—metaphor→Y* is essentially a bridge that has been established by either of the triangulation or squaring rules, as applied from within Sapper's default constructor space. When the generic concepts or schemas underlying this construction are predominantly of a spatial/perceptual nature, it becomes useful to speak of the schema *X—resemble→Y*, though computationally, both arise from the same mechanisms in constructor space. We can thus speak of these schemas as the conceptual counterparts of the computational rules that give rise to them. In this section then, we elaborate on the types of unifying schemas that an agent would need to understand the complex examples of section 2, and from this elaboration, infer the type of constructors that are computationally implicated in the comprehension process.

We adopt a compositional approach to this elaboration, arguing that the basic *—metaphor→* and *—resemble→* schemas provide sufficient building blocks to model the pragmatic phenomena of blend recruitment (both internal and external), recasting and doublethink. For instance, the basis of the *Yamaichi:Tamagotchi* metaphor of Figure 1 can be explained using the following composition of metaphor and resemblance schemas:

> *Yamaichi—metaphor→PiggyBank—resemble→**Pig***
>
> ***Pig**←metaphor—Puppy←resemble— TamagotchiPuppy*

while the visual mapping of Anatoly Chubais to a bowling pin in Figure 2 can equally be explained via the chain:

> *Chubais—metaphor→Yeltsin—resemble→**YeltsinRussianDoll***
>
> ***YeltsinRussianDoll**←resemble—BowlingPin#1←resemble—BowlingPin#2*

Our initial exploration in the domain of political cartoons show these reasoning chains, each of which is a four-fold composite of the basic metaphor and resemblance schemas, to be as complex as one is likely to find in this or other domains. Conceptually then, we can imagine a cognitive agent to have access to a set of mapping schemas that correspond to the full range of permutations of such a reasoning chain. Computationally, we can view the constructor space of a blend to be populated with the equivalent computational forms, namely, chained compositions of the triangulation rule. In algorithmic terms, this means that Sapper attempts to reconcile counterpart elements of the input spaces not simply via a single bridge, or triangulation, but via a linked chain of such bridges. There is a certain symmetry to this perspective, for as it stands, Sapper is a path-based architecture, one which operates by finding the longest chains of entities and relations in each input space that can be systematically placed in alignment. To model the pragmatic pressure of blend recruitment, we simply generalise the need to ground the conceptual end-points of these chains in a single bridge relation, to allow these end-points to be grounded at either end of a *chain* of such bridges. The intermediate concepts through which this chain is threaded act not only as computational *stepping stones*, but pragmatic mediators, for the overall metaphor. This account then is all the more cognitively plausible since it equates pragmatic pressure with computational pressure. The pragmatic phenomena we have discussed are not merely a source of difficult problems that must be addressed for the sake of completeness, but a source of efficient and tractable algorithmic constraints.

But just as there are effective cognitive limitations on the number of elements one can store in working memory, or nest in a centre-embedded clause, it is reasonable to assume that the amount of *'bridging distance'* tolerated by the metaphor faculty is similarly bounded for reasons of computational tractability. Sapper currently operates with a maximal chain size of four bridge schemas or triangulations, but again, this proves effective for even the most complex metaphors and blends we have encountered so far. It remains to be seen whether the computational limit is pragmatically determined – that is, whether the context dictates how much computational effort should be applied. For instance, one might expect that political cartoons demand more cognitive expenditure than, say, advertising images. This conjecture is the subject of current on-going research.

5. Computational Accounts of Similarity

Intrinsic to our pragmatic motivation for, and computational account of, intermediate blend recruitment is the notion that semantic similarity is more than an issue of pure structure-matching. Rather, both structural isomorphism (cf. Gentner and Markman, 1994; Markman and Gentner, 1996) and attribute-based measures of similarity (cf. Tversky 1977) must be combined to provide a complete picture of the cognitive processes involved in complex, blend-centred metaphors. This observation is not a new one: models such as ACME and LISA already embody this pluralist view of similarity in a computational framework, the former by hard-wiring similarity biases between mapping nodes in its constraint network, and the latter by gauging the extent to which two input structures are grounded in the same set of semantic features.

But Sapper stakes a greater theoretical investment in this claim, arguing that structural isomorphism is not sufficient in itself to secure an analogical or metaphor interpretation. While ACME and SME, and to a lesser extent LISA, will produce interpretations that are not grounded in attributive, or literal similarity (or what SME tellingly denotes *'mere similarity'*), Sapper nevertheless insists that all of its interpretations are, at some level, cashed out in terms of mappings that exhibit an immediate *sameness*. This sameness may be exhibited via a direct sharing of attributes (e.g., *Pen* and *Sword*, which serve to ground *Author* as *General*, are both *Long*, *Narrow*, *Sharp* and *Hand-Held*), or through the recruitment of intermediate blends that share different attributes with each end of the mapping. In category terms, the former corresponds to a simple case of direct similarity, while the latter forms the basis for a *family resemblance* (see Lakoff and Johnson, 1980; Lakoff, 1987).

5.1. Similarity as a Strategic Bias

The particular mix of attributive and structural similarity that an agent employs does not seem to be entirely deterministic, as different agents may exhibit biases toward one or the other. For example, Hofstadter *et al*. (1995) note that some agents demonstrate a strong emphasis for literal similarity when solving analogical problems, while others more fully embrace the analogical, or non-literal potential of a situation. Similar observations regarding the relative distribution of literal versus analogical modes of problem-solving are presented in Falkenhainer, Forbus and Gentner (1989). One should therefore expect a cognitively-plausible computational model to exhibit similar biases, for though such a model may be parameterised to reflect the continuum of possible behaviours, as is Hofstadter *et al.*'s Copycat architecture, a coherent model must champion just one bias at a time.

As described in this paper, Sapper exhibits a certain *'smart-aleck'* intellectualism in its balance of attributive and structural similarity. In attempting the find the longest possible semantic chains that are isomorphic (or near-isomorphic, with slippage) across domains, Sapper is in fact attempting to construct the most structurally elaborate, yet semantically defensible, theory of how two domains can be analogically reconciled. For instance, consider again the Tabletop scenario of Figure 4, in which Henry's action is to touch his coffee-cup. The most obvious, and literal, analog to this action would have Eliza touch her coffee-cup in turn, which in Sapper terms corresponds to an alignment of the shortest possible semantic chains (namely, *Henry—touch→Cup#1—isa→CoffeeCup* and *Eliza—touch→Cup#2—isa→CoffeeCup*). In contrast, Sapper determines that the most elaborate analog would instead have Eliza touch her plate of chocolate cake, since the coffee cup to her right is analogous to the coffee pot to Henry's right, while the fork and glass to her left are analogous, in that order, to the knife and glass on Henry's left. Structural similarity thus captures the geometric intuition of this scenario. But in addition, there are strong semantic (i.e., attributive) reasons for matching a plate of chocolate cake to a cup of coffee: one lies on a plate, while the other sits on a saucer; both contain dark food-stuffs, each of which may contain cream; one contains cocoa, while the other contains caffeine, both of which are addictive luxury items originating in the third world. Indeed, Sapper would readily apply a similar 'smart-aleck' rationale to reconcile the plate of chocolate cake and fork with a mirror containing a heap of cocaine and a razor blade, should such paraphernalia ever become a fixture of TableTop's universe.

Reactions are typically mixed to this form of 'clever', over-produced analysis, again suggesting that context may be the ultimate arbiter of what kinds of similarity and mapping strategies one should use to integrate two conceptual spaces. Nonetheless, though Sapper's elaborate and structurally-deep mappings may not always be considered apt, the existence of a separate constructor space does provide the theory with the necessary growing room to become even more goal-oriented and pragmatically-sensitive in the future.

5.2. The Importance of Attribute Grounding in Metaphor

Each of the examples considered so far are evidence for the importance of *'superficial appearance'*, or perceptual attribution, in the determination of metaphoric aptness. For in each case we see that deep structural and relational isomorphism is insufficient by itself to make a mapping pragmatically coherent. For instance, while there are numerous causal reasons for viewing the Yamaichi brokerage as a Tamagotchi puppy, the metaphor works ultimately because of an ingenious choice of visual carrier – a piggy bank – to bridge the imaginistic void between a financial institution and an electronic 'animal'. As argued in Tourangeau and Sternberg (1981), the attributes of the tenor and vehicle concepts, whether deep or superficial, establish the conceptual distance that the metaphor must bridge, and aptness is in turn measured with regards to this distance. The examples of Figures 1, 2 and 3 support the view that attributes as well as relations are vital to the aesthetic and pragmatic evaluation of a metaphor.

Such a view is contrary to the accepted wisdom of analogical research. Consider for instance that Gentner (1983) defines analogy as a structural process in which relations, rather than attributes, are the key players; attributes, represented as one-place predications, are effectively side-lined in such a view of analogy (see Langley and Jones 1988 for a critique of this approach). Later, Falkenhainer, Forbus and Gentner (1989) entrench this belief in the SME computational model by allowing attributes to affect the outcome of a mapping only when they are causally required by a higher-level relational structure. But as is clear from the examples of Figures 1,2 and 3, the attributes that appear to make these metaphors work so well do not fall under this category of 'causal requirement'. The causal mechanics of the Tamagotchi domain, for instance, centre around a central agonist—a virtual puppy—whole virtual life and death is controlled by an antagonist—the player—who provides food and sleep at various intervals. These mechanics are mirrored by those of financial regulation, under which a government can likewise affect the solvency (life) and bankruptcy (death) of a financial institution by providing money (food) or trade-suspension (sleep) at timely intervals. The perceptual properties of banks or puppies have no sway in the structural fit of these domains.

The structural mapping underlying a metaphor essentially provides a *causal theory* of why one domain, the tenor, can be viewed as another, the vehicle. While it is a responsibility of a metaphor to construct a compelling and coherent causal theory, a metaphor is also responsible for *communicating* this theory effectively. These complementary pressures of competence and performance also apply to predictive analogies, since one cannot feel safe in transferring knowledge from one domain to another if one feels the domains are not equivalent in some pragmatic sense. So metaphors (and analogies) are subject to two main communicative goals: firstly, a metaphor must convincingly equate the elements of one domain with those of another; and secondly, given this cross-domain correspondence as a basis, allow relational elements of the vehicle domain to be transferred to that of the tenor. This second goal crucially depends on the effectiveness of the first. The effect of the visual Tamagotchi metaphor is to say *'financial institutions are much more like government pets than you might have thought'*, and it is the establishment of this identity that allows the metaphor to convey its content.

5.3. Domain Incongruence

By using the triangulation rule to ground high-level structure in low-level semantic attribution, Sapper is able to model on important aspect of metaphor interpretation that has been largely ignored in the literature of *classical* structure-mapping models — that of *Domain Incongruence* (see Ortony, 1979; Tourangeau and Sternberg, 1981). While the classical view of an attribute is of a flat feature or single-arity, unnested predication, the same attribute can possess different meanings in different domains, and it is this plurality of meaning that serves to ground a metaphor between two domains. For instance, when one claims that a *'tie is too loud'*, one is using the attribute *Loud* in both an acoustic and a visual sense; a *Loud* tie is a *Garish* tie whose colours invoke a visual counterpart of the physical unease associated with loud, clamorous noises. But for *Loud* to be seen as a metaphor for *Garish*, such attributes must possess an internal semantic structure to facilitate the mapping between both. That is, attributes may possess attributes of their own (e.g., both *Loud* and *Garish* may be grounded in terms of the attributes *Sensory*, *Intense* and *Uncomfortable*). The division between structure and attribution is not as clean a break then as models such as SME, ACME, and to a lesser extent, LISA predict. Each of these models is predicated on the notion that semantic features are ontologically and representationally distinct from conceptual relations; even LISA, which is a significant departure from ACME in its desire to ground structure in a generic space of semantic features, assumes that these features are the terminal nodes of a representation, primitive atoms of meaning that do not possess internal semantic structure of their own. Yet, as our examples throughout this paper demonstrate, relational structure often *blends into* feature attribution, so both should be handled homogeneously.

5.4. Psychological Essentialism

In the words of Medin and Ortony (1989), a form of *'psychological essentialism'* seems to be at work here. The psychological essentialism thesis states that even though people realise that objects do not necessarily possess a defining essence that makes those objects what they are, e.g., that a sleeping potion is what it is precisely because it possesses a *dormative* essence, people frequently act as if they do. Medin and Ortony further argue that such essences are psychological correlated with the superficial attributes of an object, that is, people often reason about the deep properties of an object or idea at a superficial level because of a belief that attributes at this level are indicative of deeper causal properties. Psychological essentialism thus provides a cognitive shortcut around having to reason with the complex causal models of an object, instead allowing people to reason with simpler, attributive descriptions. In the Tamagotchi example, the metaphor works at a deep level because of a superficial similarity between the visual representation of Yamaichi as a piggy-bank and the visual representation of the virtual pet as a puppy. This surface fit suggests that whatever causal processes can be applied to the virtual pet can also be applied to the Yamaichi brokerage, and this suggestion in turn supports the structural theory communicated by the metaphor. Indeed, the reader does not need to know what these processes are, merely that they are somehow applicable in the tenor domain. This makes the metaphor pragmatically attractive as it conveys a sense of depth even to those readers who know little about the workings of a Tamagotchi.

Psychological essentialism suggests that people may infer deep causal properties of a domain from the surface attributes of that domain, and this indeed does seem to be the case. Consider for instance the examples of Figures 4(a) and 4(b), which present two alternative visual metaphors for the representation of politicians. Figure 4(a) represents European politicians as columnar statues in a Roman edifice, while 4(b) depicts British politicians from the 1980's Thatcher government as playing cards. A Roman-style structure is an apt encompassing metaphor for the political elite of the European Economic Community, given that the Roman empire is itself a good metaphor for the EEC (which of course was founded with the *Treaty of Rome*). Each of the depicted politicians thus acts as a pillar of the European community, but is represented visually using the blend of a Roman statue, which conveniently combines the surface features of the politicians concerned (e.g., *Male, Fat, Pompous*, etc.) with the necessary structural properties of a pillar (e.g., *Strong, Stone, Foundation, Support*, etc.).

Figure 4: Visual Blends in which politicians are viewed iconically as Roman statues (source: The Economist, December 10th 1983) and alternately as Royal playing cards (Source: The Economist, August 22nd 1981); again the pragmatic context dictates which view is the most appropriate.

The transformational chain linking the concept *HelmutKohl* to *Pillar* then is arguably *HelmutKohl—resemble→RomanStatue—resemble→Pillar*. Additionally, relations such as *Pillar—support→Structure* can be transferred into the tenor domain to imply that Kohl and Mitterand each support and maintain the structural integrity of the EEC, while the lack of structural support offered by the Thatcher representation is evidence of some discord in the EEC domain.

Interestingly, the Thatcher statue is depicted running from the scene in dereliction of her columnar role, yet ‘running’ is a predication that applies to the tenor domain rather than that of the vehicle. So it seems that certain attributes of the tenor domain (such as Animacy) are literally projected into the vehicle domain, in violation of the perceived unidirectionality of metaphor (see for instance Lakoff and Johnson 1980a,b) and then *re-projected* in an analogical fashion back into the tenor domain to imply a desertion of duty. This projection/re-projection is arguably facilitated by a combination of blending theory and psychological essentialism; indeed, we suggest that blend theory is the perfect vehicle in which to study studying psychological essentialism, since blends allow surface and causal properties from two domains to be combined into a mediating conceptual model, thus side-stepping the thorny issue of directionality. In this instance, because Madame Thatcher shares so many perceptual attributes with her marble counterpart, it is easy to assume that they also share deeper causal properties that are correlated with these attributes (e.g., the presence of legs implies a running ability, even if these legs are made of stone). Interestingly, this form of surface→causal correlation pervades the logic of children’s stories, where anthropomorphism is facilitated by the presence of superficial similarities (e.g., gingerbread-men dance due to the presence of gingerbread arms and legs, while animals talk due to the presence of mouths, and so on).

The example of Figure 4(b), which depicts British ‘Tory’ politicians as playing-cards undergoing a ‘cabinet re-shuffle’, is interesting because it employs a visual blend (*Royal-Card* = *Royalty* and *Card*) not in a mediating role but as an end-point of the metaphor. Royal picture-cards are used to depict cabinet ministers of the period, such as Geoffrey Howe and Michael Heseltine, while undecorated numbered cards are used for non-cabinet ministers of the back-benches. The concept of a Royal-Personage thus acts as a mediator in the mapping, the implication being that ministers with cabinet portfolios are the aristocrats of government who lord it over the commoners (i.e., the faceless number cards). The fact that one number-card *is* decorated with a face suggests that a back-bencher is likely to be promoted to royal cabinet status. Individual cabinet ministers are mapped to specific Royal figures then: Geoffrey Howe, for instance, is mapped to a Queen, suggesting perhaps that he possesses certain feminine (i.e., soft) qualities, while the younger Michael Heseltine is mapped to the Knave, suggesting qualities such as cunning, rakishness and mischief. Attributes such as Power and Age are major determiners of the *Politician-as-Royal* mappings, illustrating the importance of attribute mapping in metaphoric aptness.

Interestingly, the domain of Royal personages is richer in control and dominance structures than that of card games alone. Individual cards *do* have an associated numeric ‘ranking’ and ‘status’ of their own, but the relationship amongst cards (which trumps what, and so on) is largely game-dependent. In contrast, real Kings always outrank Queens, which outrank Knaves (knights), all of which outrank and control commoners. Much of the message of the metaphor is carried then by the conceptual structure of the mediating royal figures, rather than the elements of the vehicle domain itself. Indeed, the overlap between the surface features of a royal personage and its playing-card depiction suggests that they share the same psychological essence, i.e., one reasons about playing cards in a game-neutral way by reasoning about their royal counterparts.

5.5. Congruent Verbal Metaphors

One might argue that the visual metaphors we have explored are cognitively and pragmatically distinct from the verbal metaphors one encounters in a text, since the graphic requirements of the former give greater prominence to the surface attributes of the concepts involved. However, since metaphor is a product of the imagination, one expects every metaphor, linguistic, graphical or otherwise, to 'paint a picture' in imaginistic space. Most textual metaphors thus have visual counterparts, and vice versa, and the pragmatic effectiveness of these metaphors depends on the coherence of the underlying conceptual picture.

Indeed, this connection between words and pictures is evident in the ability of a visual metaphor to evoke established linguistic metaphors. In this respect the metaphors of Figures 1 through 4 are but graphical depictions—albeit highly polished and elaborated depictions—of existing figures of *speech*. Intuitively, two input spaces will seem all the more similar if there exists a set of verbal metaphors (or clichés) with which they are congruent. For example, the piggy-bank of Figure 1 is shown in a *'belly-up'* position to communicate a message via a conventional corporate and business metaphor, *'to go belly-up'*. Likewise, the bowling scene of Figure 2 is congruent with a variety of established political metaphors, such as *'fall from grace'* and *'to be rocked by scandal'*. The circus visuals of Figure 3(a) are congruent with the linguistic metaphor *'to jump through hoops'*, while Figure 3(b) is a rather direct visual interpretation of the figure of speech *'to have a fling'*. In Figure 4(a) we see Helmut Kohl and his companions act as *'pillars of the community'*, where in this case the community is that of the EEC. And in Figure 4(b) we see Margaret Thatcher perform a *'cabinet shuffle'*, a long-established linguistic metaphor which is itself congruent in this context with the idea of a figurative *'king-maker'*. Other linguistic metaphors, such as *'to have an ace up one's sleeve'*, or *'to deal from the bottom of the deck'*, might also be applicable in this context. These linguistic metaphors, which are common currency in their respective domains of discourse, resonate with, and serve to reinforce, the central message of the visual depiction.

This in turn is further evidence that the pragmatic effectiveness of a metaphor cannot be judged wholly at a deep conceptual level. For not only must a cognitive agent consider the coherence of the surface attributes involved, but the typical manner in which a metaphor's message is conveyed at the linguistic surface, if its full pragmatic force is to be appreciated. It follows that if a wholly conceptual and non-linguistic theory of metaphor, such as SME, ACME or Sapper, is to recognise any structural congruences with established linguistic metaphors, it will require structural *traces* of these metaphors to be evident at the conceptual level. These traces, which indicate the conceptual underpinnings of the metaphors concerned, will allow the conceptual model to indirectly hook into the surface form of those metaphors; that is, if a conceptual interpretation of a visual metaphor is structurally coherent with such traces, then for the most part it should also be coherent with their linguistic expressions.

Neither SME, ACME or LISA currently recognise the need for metaphoric traces; neither do they provide the conceptual machinery to do so, since a trace implies a memory of past metaphoric reasoning, and neither model employs an explicit model of memory (Langley and Jones (1988) echo this criticism). However, because Sapper lays down conceptual bridges between entities when acquiring new concepts and interpreting novel metaphors, these bridges *are* effective traces of its past reasoning. Underlying many clichéd linguistic metaphors then, such as *'the pen is mightier than the sword'*, will lie a series of one of more bridges, such as *Pen : Sword*. For instance, *Author* as *General* can connect to this metaphoric bridge via the isomorphic relational chains *Author—control→Pen* and *General—control→Sword*. The extent to which a new metaphor, visual or otherwise, resonates with existing metaphors in memory can thus be measured, by counting how many existing conceptual bridges the current metaphor can systematically be connected with. In effect, the mechanics are the same as those discussed for metaphor recruitment in section two. Strictly speaking, congruence and recruitment are separate cognitive phenomena, since the former is an optional aid to pragmatic coherence while the latter is a necessary step in composing an interpretation. At the computational level, however, each can be treated identically in terms of structural isomorphism.

5.6. Similarity as a Tractable Guide to Unpacking

Recall that unpacking is a crucial facet of blend usage that allows a blended space to be *'deconstructed'* to reveal its conceptual origins. As such, this ability to unpack underlies the introspective processes by which people reconstruct the basis of a conceptual integration, as when, for instance, one revisits the conceptual metaphor *SpaceTime as a Fabric* to reconsider and better understand the *BlackHole* blend. But not all blends are so easily excavated; people are not so quick to unpack the *HotDog* blend into its original spaces *WeinerWurst* (a Viennese sausage, similar to a *Frankfurter*) and *Dachshund* (or *SausageDog*). To facilitate unpacking, the relations that link into a blended space must be readily navigable in reverse, meaning that such links should be few in number or easily discriminated in terms of their salience. When this is not the case, the process of unpacking becomes a highly ambiguous search process of intractable proportions, prompting the cognitive agent to simply give up. But our sense of what constitutes semantic similarity can help to prune this search process, and heuristically suggest what relations are worth retracing into other spaces.

Consider the concept *CelticTiger*, a blend which describes the currently booming Irish economy and its thriving computer industry. This blend is an obvious extension of the established concept *AsianTiger*, which in turn blends the concept *Asia* with the metaphor *AgressiveEconomy as Tiger*. Note that the concept *Tiger* is carried directly into the new blend, despite the obvious fact that Ireland is not at all associated with tigers. From a relational perspective, the best counterpart one can imagine to an Irish tiger is the concept *Wolfhound*, since wolfhounds and tigers are fierce, aggressive, proud and mythological in stature. Nonetheless, the blends *CelticWolfhound* and *WolfhoundEconomy* simply do not work. In contrast, economists describe Tunisia not as an *AfricanTiger* but as an *AfricanLion*. Why is this?

This blend may be more successful because it facilitates rather than hinders the unpacking process: to many people's minds, Africa and Asia are economically and politically similar, while lions and tigers are frequently conceptualised as mere *regional variants* of each other. This suggests that one can reach the metaphor *AgressiveEconomy as Tiger* from the concept *Lion*, but not from the concept *Wolfhound*. We argue that the latter route is mentally pruned during the unpacking process simply because it does not seem a promising route, for despite occupying an equivalent structural position in the Celtic culture, it simply is not similar enough to our concept of Tiger to be considered a tractable option. Psychological essentialism, which here claims that the metaphoric essence of one concept (*Tiger*) can be transferred to another concept which is superficially similar (*Lion*), may thus observe certain similarity thresholds in its operation. It remains for us to determine, and algorithmically specify, these thresholds.

6. Conclusions

Metaphor is a phenomenon that defies neat cognitive compartmentalisation, pervading as it does not only linguistic, but visual, phonetic and physical realms. For such reasons metaphor has become one of the dominant research focii in Cognitive Linguistics, a field of endeavour which denies that language can ever be placed into its own, clinically separated cognitive module. The examples considered in this paper clearly illustrate this cross-module nature of metaphor, while providing further cognitive support for several theories of concept integration, in particular, the conceptual blending theory of Fauconnier and Turner (1994, 1998) and the Sapper theory of structural metaphor as posited by Veale *et al.* (1994, 1995, 1996, 1997). The latter is a computational framework that additionally demonstrates how the phenomena we discuss can be translated into, and considered from, an algorithmic perspective.

This algorithmic perspective is vital if pragmatic pressures are to be handled in an ecologically valid fashion: that is, pragmatic pressures on interpretation must reflect underlying computational pressures on the tractability of this interpretation, if they are to be properly conceived as *pressures* (rather than mere curiosities) in the first place. Our computational account respects this dependency between levels, by showing how blend recruitment not only enhances the communicative force of a metaphor/blend, but how it can also guide the computational model in its efforts to efficiently integrate different conceptual spaces in a semantically grounded fashion.

We thus argue that the need to rationalise this integration of spaces in terms of semantic similarity is both cognitively and computationally attractive. It is cognitively attractive because it reflects people's use of metaphors and blends which are themselves complexes of other, congruent metaphors, and thus explains how novel combinations do not need to be constructed from whole cloth, but can be elaborated from an established metaphoric ground. Metaphors are used to convey a world-view, to communicate to an audience why a certain perspective is valid and compelling. As such, metaphors are arguments, and a compelling argument is one whose main points (or components) are easy to establish in their own right. It follows that complex metaphors facilitate cognitive digestion when their parts are themselves elements of a much-traded and readily accepted metaphoric currency. One's sense of semantic similarity, and the ability to reconcile concepts in terms of their intrinsic properties, provide the basic *exchange mechanism* for this currency.

Psychological essentialism plays a key role in this exchange, allowing one to infer deep properties on the basis of surface similarities. Essentialism often provides the rationale for blend recruitment, allowing a concept to act as an intermediary because it shares superficial properties with one end of a mapping, and deep properties with the other. Consequently, one performs an act of essentialism when transferring the deeper properties of one side of the mapping to the other (e.g., as in believing that Yamaichi, a brokerage, can actually die and go 'belly up'). Indeed, the notion that people are often slaves to this form of shallow reasoning is captured by the popular idiom *'never judge a book by its cover'*.

But again, psychological essentialism is more than a cognitive curiosity, for it greatly enhances the tractability of conceptual integration. By allowing integration to proceed on the basis of attributive, rather than deep relational similarity, the complexity of the mapping process is much simplified. As demonstrated in the case of the Sapper model, existing conceptual bridges (dormant or active) can be exploited to this end, without the need to explore the full relational consequences of a recruited mapping. Recall that such recruitment is intended to facilitate a higher-level mapping, not hinder it, yet this is exactly what would happen if an agent was forced to consider every structural ramification of every low-level mapping. Psychological essentialism yields a computational short-cut, a safety valve on the tractability of the interpretation process.

We conclude on this theme of computational felicity, by noting that the model of blend recruitment presented in this paper may also shed a useful computational perspective on another intriguing aspect of Fauconnier and Turner's theory of blending, namely the *metonymy projection principle*. Since metaphors and blends typically serve the communicative purpose of throwing certain elements of a domain into bas relief, while de-emphasising others (e.g., see Ortony, 1979), this strengthening of associations frequently causes the relational distance between the tenor and its highlighted association to be fore-shortened in any resulting conceptual product.

Fauconnier and Turner cite as an example of this principle the concept *GrimReaper*, a blend which metaphorically combines the concepts *Farmer* and *Death*. In the latter domain, the concepts *Skeleton* and *RottingClothes* are causally associated with *Death*, via the intermediate concepts *Decompose, Rot, Coffin, Funeral, Graveyard*, and so on. But in the resultant blend space, *Skeleton* and *RottingClothes* become directly associated with *Death*, and are used together as an explicit visual metonym; the Grim Reaper is thus conventionally portrayed as a scythe-carrying skeleton, wrapped in a decrepit cloak and cowl. We see a similar instance of this phenomenon in the Tamagotchi example of Figure 1, in which the associations between *Yamaichi*, a rather lofty brokerage, and the concepts of *PersonalSavings* and *SmallInvestor* are strengthened by the use of a *PiggyBank* as a visual metonym. This has the effect of personalising the metaphor and making its consequences more relevant to the intended audience, the bulk of which will themselves be small, rather than corporate, investors. In both these cases, metonymic short-cuts emerge because an intermediate blend is recruited that provides a shorter path to the relevant associations. *Skeleton* serves as a rich visual analog of *Farmer* (both have arms, legs, torso, head, etc.) while evoking certain abstract properties of *Death*, whereas *PiggyBank* is a rich visual analog of a *TamagotchiPuppy*, while sharing key abstract properties with *Yamaichi*.

The computational account we provide to explain this phenomenon of blend recruitment may thus provide an algorithmic basis for much of what passes for *metonymic projection*. It remains as a goal of future research to establish other aspects of conceptual integration that can neatly be accommodated within this general computational framework.

7. References

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

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.

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

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.

French, R. (1995). *The Subtlety of Sameness*. Cambridge: MIT Press.

Gentner, D. (1983). Structure Mapping: A Theoretical Framework for Analogy. *Cognitive Science* **7**(2), pp 155-170.

Gentner, D. and Markman, A. (1994). Similarity is like analogy. In C. Cacciari (Ed.), *Proceedings of the Workshop on Similarity at the University of San Marino*. Milan, Italy: Bompianai.

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**, pp 295-355.

Lakoff, G. and M. Johnson. (1980a). The metaphorical structure of the human conceptual system, *Cognitive Science* **4**, 195 - 208.

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

Langley, P. and R. Jones. (1988). A Computational Model of Scientific Insight, in *The Nature of Creativity*, R. J. Sternberg (Ed.). New York: Cambridge University Press.

Medin, D. and A. Ortony. (1989). Psychological Essentialism, *in Similarity and Analogical Reasoning*, Vosniadou, S. Ortony A. (Eds.) Cambridge University Press, 1989.

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

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

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

Turner, M. and G. Fauconnier. (1995). Conceptual Integration and Formal Expression, Journal of Metaphor and Symbolic Activity **10**(3).

Markman, A. B., and Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology, 25*(4), 431-467.

Tversky, A. (1977). Features of similarity. *Psychological Review*, **84(**4), 327-352.

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

Veale, T, D O'Donoghue and M. T. Keane. (1995). Epistemological Issues in Metaphor Comprehension: A Comparative Analysis of Three Models of Metaphor Interpretation, presented at *ICLA'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, *Cognitive Linguistics: Cultural, Psychological and Typological Issues (forthcoming).*

Veale, T. and M. T. Keane. (1997). The Competence of Sub-Optimal Structure Mapping on 'Hard' Analogies, *in the proceedings of IJCAI'97, the Fifteenth International Joint Conference on Artificial Intelligence, Nagoya, Japan. San Mateo, CA: Morgan Kaufmann.*