Belief Modelling, Intentionality and Perlocution in Metaphor Comprehension

Tony Veale
Hitachi Dublin Laboratory &
University of Dublin, Trinity College,
Dublin 2, IRELAND
tony.veale@hdl.ie

Mark T. Keane
Dept. of Computer Science,
University of Dublin, Trinity College,
Dublin 2, IRELAND
mark.keane@cs.tcd.ie

Abstract

Metaphor is an elegant, concise, often startling communicative form which is employed by a speaker as a means of conveying a state of affairs to a hearer; as such, it deserves to be analysed as a speech-act, with a particular illocutionary intent and perlocutionary effect. This paper describes a hybrid symbolic/connectionist model of metaphor (SAPPER by Veale & Keane, 1993), which incorporates elements of the belief ascription model of (Wilks, Barnden & Wang, 1991). This extended framework provides a suitable computational environment for analysing the illocutionary intent of the speaker, and perlocutionary effect upon the hearer's belief space, of a broad class of metaphors with an observable ameliorative/pejorative connotation.

1. Introduction

The view, promoted by Lakoff & Johnson (1980), that metaphor plays a fundamental structural role in organizing our conceptual systems, rather than being a deviant rhetorical effect, is now generally accepted. Indeed, this idea drives many recent computational models (Weiner, 1984; Martin, 1990; Way, 1991;Veale & Keane, 1992a, 1992b). Such analyses are based on the de-contextualised "contents" of the metaphor and ignore the communicative context in which it is uttered. However, Davidson (1979) has warned against confusing the content of a metaphor with the intended effect of a metaphor. In his view metaphor, is not a *carrier* of meaning, but a *recipe* from which meaning is constructed. The real meaning of a metaphor lies in the changes it makes to belief structures of the reader.

So, in order to capture the richness of their metaphoric utterances they need to be viewed as speech acts (see Austin 1962). From this perspective the metaphor is uttered by a speaker with a specific communicative intent, in the context of a *speaker-specific* world model, and subsequently interpreted by a listener relative to a local world model which is similarly *listener-specific*. To arrive at a full interpretation of the utterance, a system must therefore characterise that component of meaning which is common to the world models of speaker and listener, before the *information content*, as opposed to the semantic content, of the metaphor can be represented. So, a proper analysis requires the system to model the belief structures of the speaker and listener relative to each other in the context of the tenor and vehicle concepts. Recognising the communicative intent of the speaker is also important if the speaker's beliefs are to be correctly modelled by the listener (e.g., is the speaker conveying a pejorative or ameliorative account of the tenor?).

Various researchers, notably Ballim, Wilks & Barden (1991), have subsumed the metaphoric process within a larger framework of belief ascription (see also Wilks, Barnden & Wang, 1991). Wilks et al. treat metaphor as a structural/propositional transference between domains and show how some properties of metaphor naturally arise out of a model of belief space *amalgamation*. The work of Wilks et al. informs much of what follows, but we are concerned with various other properties of metaphor which fall outside their framework; properties such as *domain incongruence* (see Tourangeau & Sternberg, 1981) and *domain interaction* (see Richards, 1936; Black, 1962), which seem to transcend simple transference of predicational structure between domains.

So, our emphasis is more on the role of belief ascription and speech act analysis in metaphor interpretation, rather than on belief ascription in itself. In sections 2 and 3, we present a memory model for capturing aspects of belief ascription and metaphor interpretation. In section 4, we demonstrate how the activation dynamics of the Sapper connectionist network can be employed to derive a quantitative measure of both the emotive and communicative force of a metaphor, while sections 5 and 6 illustrate how various heuristic rules of belief revision capitalise upon this measure. The paper ends with a summary and conclusions in section 7.

2. Overview of the SAPPER Framework

The model of metaphor and belief presented in this paper is best elucidated within the computational framework of the Sapper network system (see Veale & Keane, 1993). The Sapper framework is a hybrid symbolic/connectionist model, which posits metaphor interpretation as a dual process involving the building of symbolic structures in a semantic network, and the propagation of activation across this network. In Sapper, metaphor comprehension involves the construction (or more accurately, the *awakening*) of new cross-domain linkages, which serve as bridges to bind the *analogue-pairs* established by the metaphor. Metaphor novelty is reflected by the amount of structure added to the network. In short, metaphor is modelled as a dynamic, constructive, conceptual phenomenon, which evokes a response in a reactive fashion from an adaptive and accommodating knowledge-base.

A Sapper network is a localist graph in which nodes represent concepts, and where arcs represent semantic relations

between these concept nodes. The symbolic part of the system analyses the network for particular consistencies of structure, and uses various rules to modify the network. For instance, new (but dormant) connections may be constructed on the basis of consistencies found. Such dormant network linkages represent merely *plausible*, rather than fully *established*, semantic relations, and are thus not operative carriers of activation. Sapper's connectionist component controls the propagation of activation energy or *zorch* throughout the network (see Hendler 1989), when activation flows from concept nodes evoked by the metaphor (i.e., the tenor and vehicle concepts). Under suitable conditions, propagated activation can awaken the dormant connections established by the symbolic rules, transforming them into active conceptual bridges between the domains of the tenor and vehicle concepts.

The symbolic component uses two constructor rules: the triangulation rule and the squaring rule. The *triangulation rule* is invoked whenever two concept nodes share a common association or superclass, establishing a link between them. The *squaring rule* is a second-order constructor that acts on the links built by the triangulation rule, building bridges upon bridges, each new linkage extending the inter-domain *reach* of the last. Consider how Sapper deals with the "surgeons are butchers" metaphor (see Figure 1). The triangulation rule notes common associations between aspects of the surgeon and butcher concepts (e.g., MEAT:HUMAN-FLESH, SLAUGHTER & DEATH, INTERNSHIP:APPRENTICESHIP & EXPERIENCE) and lays down dormant linkages between the schemata BUTCHER and SURGEON, HUMAN-FLESH and MEAT, CLEAVER and SCALPEL, APPRENTICESHIP and INTERNSHIP, and SLAUGHTER and SURGERY. The squaring rule then *reinforces* the square-structured bridgework SURGEON:BUTCHER & MEAT:HUMAN-FLESH, and SURGERY:SLAUGHTER & SCALPEL:CLEAVER.

Following the construction stage, activation is propagated from the matriarch nodes of the metaphor, SURGEON (tenor) and BUTCHER (vehicle), causing the dormant linkages between SURGEON:BUTCHER, MEAT:HUMAN-FLESH, SURGERY:SLAUGHTER, SCALPEL:CLEAVER, and INTERNSHIP: APPRENTICESHIP to be awakened as full, activation-carrying bridges. A dormant linkage is awakened in Sapper whenever it serves as channel for competitive activation waves originating at different matriarch nodes: that is, whenever it is seen as constituting a cross-domain bridge between the domains of tenor and vehicle. The opening of these bridges allows activation to flow freely between the tenor and vehicle domains, altering the activation dynamics of the network in such a way that the tenor *actually* interacts with the vehicle at a conceptual level. The activation patterns of BUTCHER, SLAUGHTER, MEAT and CLEAVER interact with those of SURGEON, SURGERY, HUMAN-FLESH and SCALPEL to produce a response to the metaphor. So, a Surgeon is seen, through the lens of the metaphor, to be an altogether less skilful and precise tradesman, performing surgery which is akin to the slaughter of innocents, while amidst the pain and screams of fear, wielding a blood-stained scalpel to slash and chop liberally into human *meat*. The metaphor operates both ways, denigrating surgeons, but elevating butchers, who are now seen to be that much less clumsy, careless and imprecise, and altogether more professional. This bi-directional (though not symmetric) interplay of different conceptual domains captures the interaction view of Richards (1936) and Black (1962), in a way which the simple transference of predicational structure, as advocated by Wilks, Barnden & Wang (1991), in their ViewGen model, does not.

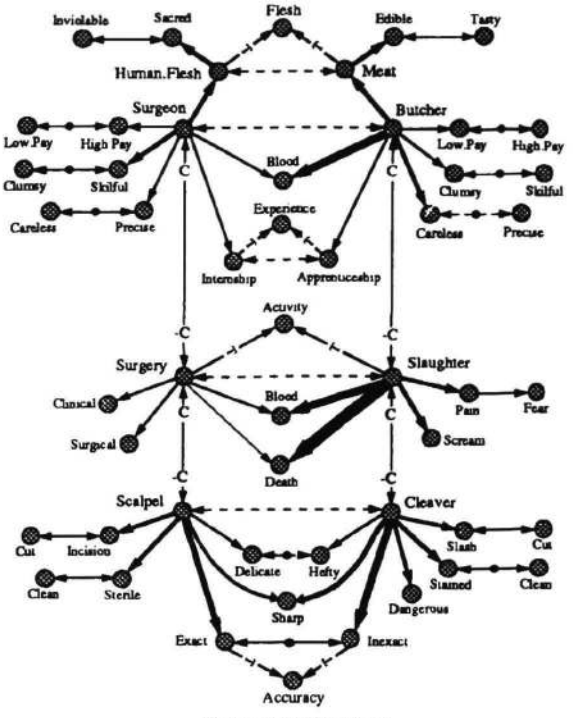

Figure 1: The Sapper network structure underlying the metaphor "Surgeons are Butchers". (Key: Dashed bi-directional arrows depict dormant **Family Resemblance** relations, while unbroken unlabelled bi-directional arrows depict established family resemblances; unlabelled uni-directional arrows depict **Attributive / Associational** relations; bi-directional arrows labelled C -C depict **Control** relations; unidirectional arrows labelled T depict **Taxonomic** relations; and • depicts an **inverter**, which transforms inhibitory activation into excitation, and excitatory activation into inhibition; Line thickness indicates relational salience).

The Squaring rule is considered second-order as it employs links rather than nodes as the evidential basis for performing structural inference. Such second-order strategies thus support an interplay between the symbolic and connectionist components of the hybrid model, for as dormant links are awakened, they may instigate further high-level inference. In effect, to perceive two concepts as being similar, it is not simply enough to recognize the commonality in associational structure, but it is often necessary to invent new commonalities (see Finke, Ward & Smith, 1992). These newly perceived

commonalities of structure underpin the phenomenon of *Domain Incongruence* (see Tourangeau & Sternberg, 1981), and represent the *territorial expansion* of a metaphor as it builds upon other, lower-level metaphors to extend its reach between evermore diverse conceptual domains. Metaphors often employ high-level, or non-literal similarities, based on attributes which hold different meanings in different domains (see Tourangeau & Sternberg 1981). Domain incongruence is exhibited in such metaphors as "Her heart was an understocked ice-box", where refrigerators and human hearts are neither empty or cold in the same manner. Rather emptiness is itself a metaphor representing human indifference, while coldness signifies a lack of compassion.

3. Belief Ascription in Sapper

A Sapper network, which is essentially a generalised graph structure comprising labelled relational arcs/links between labelled concept nodes, may be viewed as an unstructured collection of propositions (i.e., two-place predications annotated with a numeric measure of relational salience). Viewed in this manner, we can readily superimpose a belief maintenance architecture onto the Sapper network model. The belief sets of sentient agents (e.g., people, computers) are modelled as *viewpoint environments* which contain different *topic environments*, each relating to a different concept as viewed by a particular agent (following Wilks, Barnden & Wang's, 1991, terminology). The Sapper knowledge-base thus comprises two interlinked architectures: a belief-space model composed of viewpoint and topic environments, and a localist connectionist network comprising weighted links, each of which represents a distinct proposition. Each link is annotated with one or more topic environment identifiers, as the same proposition may be held by different agents, and a corresponding credibility weighting for each topic, which represents the credence each belief agent places upon the proposition. A viewpoint is considered *primed* whenever that viewpoint is currently adopted by the system, allowing Sapper to consider a proposition from within the belief spaces of different agents; the priming of a viewpoint environment automatically causes its constituent topic environments, and the propositions contained therein, to be primed also. To support this *viewpoint-switching* capability, activation is therefore constrained by the connectionist component to occur across only those proposition links contained within a currently primed topic environment.

4. Ameliorative and Pejorative Shift

The existence of multiple viewpoints within a knowledge-base allows for considerable *meaning shift* to occur in the comprehension of metaphor, as an utterance may be interpreted differently within conflicting agent belief-spaces. This issue is of particular relevance in quantifying the overall ameliorative/pejorative content of an utterance: in the metaphor of Figure 1, for instance, an ameliorative shift occurs when considering the metaphor from the viewpoint of a butcher, while a complementary pejorative shift occurs when a surgeon viewpoint is adopted. A means of gauging this ameliorative/pejorative shift between the tenor and vehicle viewpoints is necessary if the system is to accurately determine the *cognitive, or emotive force* of a metaphor; that is, a context-independent measure of the persuasive power of the metaphor regardless of the speaker's intent. Such a measure is a principal determiner in updating the belief-space of the hearer in response to the speaker's metaphor. A connectionist means of obtaining such a measure is simply implemented within Sapper, as illustrated in Figure 2:

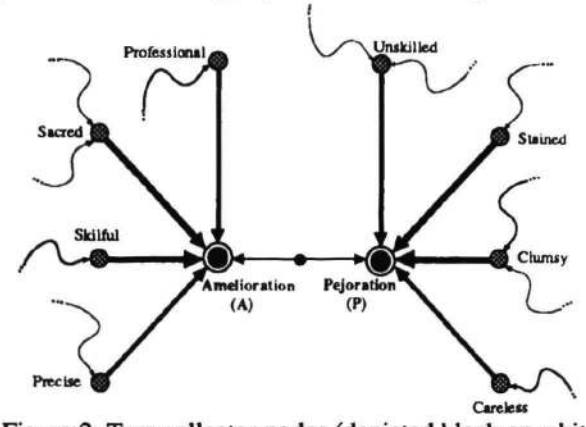

Figure 2: Two collector nodes (depicted black on white), which accept all incoming activation, but do not propagate it onwards, are used to measure the relative ameliorative/pejorative content of a concept.

As shown in Figure 2, Sapper designates two concept nodes, AMELIORATION (denoted A) and PEJORATION (denoted P), to serve as special activation collectors, which accumulate but do not redistribute activation from adjoining nodes. Descriptive concepts throughout the network which possess any measure of ameliorative/ pejorative connotation are connected to these nodes, whereby the association strength of each linkage reflects the extent of the connotation. So, by priming the viewpoint environment E of an observer O, and subsequently initiating the spread of activation, or zorch Z, from a concept node T, the activation levels of the collectors A and P will provide a measure of the relative ameliorative/pejorative view toward T by O:

- Ameliorative(O, T) = $B_O Z_T (A)$
- Pejorative(O, T) = $B_O Z_T (P)$
- View(O, T) = $B_O Z_T (A) - B_O Z_T (P)$

The zorch arriving at collector node A from node T, denoted $Z_T (A)$, in the viewpoint environment of believer O, denoted, B_O, represents a measure of the ameliorative view felt by O for T, denoted Ameliorative(E, T). The overall view (or *esteem*) held by O for T is denoted View(O, T); a positive value for this function implies O holds an ameliorative opinion of T, while a negative value suggests a pejorative opinion.

It follows therefore that T's self-esteem to defined by:

- Self-Esteem(T) = View(T, T) = $B_T Z_T (A) - B_T Z_T (P)$

The system is thus in a situation to gauge the effect upon a hearer H's self-esteem should H adopt the proposition set X (i.e., a set of one or more new propositions):

- ΔSelf-Esteem(H) = View(H+X, H) - View(H, H)
  $= B_{H+X} Z_H (A) - B_{H+X} Z_H (P)$
  $- B_H Z_H (A) + B_H Z_H (P)$

The proposition set of interest to Sapper, of course, is the set of bridges awakened in comprehending a metaphor, as Sapper is required to attach a credibility weighting to each of these new propositions in the belief space of the Hearer H. Calculation of the ensuing effects upon the self-esteem of the Hearer is therefore essential if the system is to ascertain the likelihood of the hearer adopting X, and determine the certainty / credibility rating that H will attach to X. However, it is also necessary to determine the perceived change in self-esteem experienced by the speaker S as observed by the hearer H:

- $B_H \Delta$Self-Esteem(S) = B_H [View(S+X, S) - View(S, S)]
  $= B_H [B_{S+X} Z_S (A) - B_{S+X} Z_S (P)$
  $- B_S Z_S (A) + B_S Z_S (P)]$

Thus, for a metaphoric utterance M, in which speaker S describes a tenor concept T via the vehicle concept V to a hearer H, the emotive force of M, Ef(M), is defined:

- Ef(M) = |View(H, V)| = $|B_H Z_V (A) - B_H Z_V (P)|$

i.e., the absolute ameliorative/pejorative content of the vehicle chosen to describe the tenor. However, as suggested above, the system must also consider the effect the utterance M has upon the hearer H, and the effect M has upon S as perceived by H, if the persuasive power of the metaphor is to be determined. The *Pragmatic force*, Pf(M) of a metaphor reflects this perspective shift that exists between speaker and hearer:

- Pf(M) = ΔSelf-Esteem(H) - $B_H \Delta$Self Esteem(S)

A context-dependent, or pragmatically motivated, measure of the persuasive force of the metaphor can now be formulated. This context-dependent measure is termed the *communicative force* of the metaphor, or Cf(M). If the speaker S, in uttering the metaphor M, communicates the set of propositions X to the hearer H, then Cf(M) is:

- Cf(M) $= \phi$. Ef(M) + $(1 - \phi)$. Pf(M)
  $= \phi$. Ef(M) + $(1 - \phi)$. [ΔSelf-Esteem(H)
  $- B_H \Delta$Self Esteem(S)]
  $= \phi$. |View(H, V)|
  $+ (1 - \phi).$View(H+X, H)
  $(1 - \phi).$View(H, H)
  $- (1 - \phi).B_H$ View(S+X, S)
  $+ (1 - \phi).B_H$ View(S, S)
  $= |\phi . B_H Z_V (A) - \phi . B_H Z_V (P)|$
  $+ (1 - \phi).B_{H+X} Z_H (A)$
  $- (1 - \phi).B_{H+X} Z_H (P)$
  $- (1 - \phi).B_H Z_H (A)$
  $+ (1 - \phi).B_H Z_H (P)$
  $- (1 - \phi).B_H B_{S+X} Z_S (A)$
  $+ (1 - \phi).B_H B_{S+X} Z_S (P)$
  $+ (1 - \phi).B_H B_S Z_S (A)$
  $- (1 - \phi).B_H B_S Z_S (P)$

The scaling factor ϕ represents the readiness, or *susceptibility*, of the system to accept the propositions conveyed through an utterance without recourse to the relative pragmatic position of the hearer and speaker. As defined above, the communicative force of a metaphor is thus not only dependent upon the emotive force of the utterance, but also upon the likely effect that acceptance of the metaphor is to have upon the hearer H, and upon the *openness* of the speaker S as perceived by H. If M does not relate to H in any descriptive fashion, then ΔSelf-Esteem(H) is zero and not a factor in the final interpretation. However, if M describes H positively, the communicative force is strengthened, causing H to attribute more certainty to the propositions conveyed therein, while if M describes H negatively, the communicative force is diminished and less certainty is attributed to the propositions it conveys. Similarly, if M is seen by H to describe S positively, that is, in a self-serving manner, then H is less likely to accept fully the propositions conveyed by S, while if M is seen to describe S negatively, that is, in a self-deprecating manner, H is more likely to attach greater credence to the propositional content of M. Recalling the example of Figure 1, "Surgeons are butchers", this utterance has more communicative power when uttered *by* a surgeon, or *to* a butcher, and less communicative power when uttered *to* a surgeon or *by* a butcher.

What has been achieved by this algebraic manipulation? Building upon a simple means of ascertaining the relative ameliorative/pejorative view of a concept from the perspective of a particular belief agent, Sapper is capable of quantifying the *persuasive power*, or communicative force, of a metaphor as it relates to the beliefs of the hearer.

5. Speaker Intentionality

As described and formulated in the previous section, a metaphoric utterance may effect changes in the belief space of the hearer, and some measure of the likelihood that such changes are wrought is the communicative force of the utterance. Of course, the utterance is communicated by the speaker with precisely such a persuasive goal, and following Davidson (1979), metaphors must by analysed in terms of the pragmatic goals of the speaker. When a speaker utters "My wife's cooking is a disaster", the speaker's intention is to convey a dislike of his wife's cooking. However, a propositional expression of this dislike is not contained within the metaphor itself, but arises from a pragmatic analysis of the utterance as a speech act.

Given a metaphor M uttered by S, which describes T as V, and in doing so conveys the proposition set X to H, then the following belief ascription heuristics apply:

Like/Dislike:

- When M is perceived to be pejorative toward T,
  i.e., View(H, V) < 0,
  Infer: $\Delta B_H B_S Z_T(P) \propto Cf(M)$

i.e., modify the link strengths to ensure that more zorch arrives at P from T in $B_H B_S$

- When M is perceived to be ameliorative toward T,
  i.e., View(H, V) > 0,
  Infer: $\Delta B_H B_S Z_T(A) \propto Cf(M)$

Trust/Distrust:

- When M is perceived to be pejorative toward T,
  i.e., View(H, V) < 0,
  Infer: $\Delta B_H B_S Z_T(\text{UNTRUSTWORTHY}) \propto Cf(M) \cdot \max[0, B_H B_S Z_V(\text{POWER})]$

i.e., modify the link strengths to ensure that more zorch arrives at UNTRUSTWORTHY from T in $B_H B_S$ - the extent of the modification is related to H's certainty that S holds V to be in a position of power.

- When M is perceived to be ameliorative toward T,
  i.e., View(H, V) > 0,
  Infer: $\Delta B_H B_S Z_T(\text{TRUSTWORTHY}) \propto Cf(M) \cdot \max[0, B_H B_S Z_V(\text{POWER})]$

Respect/Fear

- When M is perceived to be pejorative toward T,
  i.e., View(H, V) < 0,
  Infer: $\Delta B_H B_S Z_T(\text{FEAR}) \propto Cf(M) \cdot \max[0, B_H B_S Z_V(\text{POWER}) - B_H B_S Z_S(\text{POWER})]$

- When M is perceived to be ameliorative toward T,
  i.e., View(H, V) > 0,
  Infer: $\Delta B_H B_S Z_T(\text{RESPECT}) \propto Cf(M) \cdot \max[0, B_H B_S Z_V(\text{POWER}) - B_H B_S Z_S(\text{POWER})]$

i.e., modify the link strengths to ensure that more zorch arrives at RESPECT from T in $B_H B_S$ - the extent of the modification is related to H's certainty that S holds V to be in a position more powerful than himself/herself.

These heuristics provide a principled basis for modifying the hearer's beliefs concerning the speaker's beliefs as communicated in a metaphor. In each case, the proposed modification alters the activation dynamics of the belief viewpoint environment $B_H B_S$, in a manner which is consistent with the perceived communicative force of the utterance, as formulated in the previous section.

6. Metaphors of Mind

In dealing with metaphors which relate to a sentient tenor, that is, a tenor which is a belief agent in its own right, the system must do more than determine those beliefs the hearer will ascribe to the speaker (the perlocutive effect); it is also necessary to determine those beliefs that the hearer perceives the speaker to ascribe to the tenor. To model this ascription from speaker to tenor, Sapper employs a fundamental *metaphor of mind*, that of concept association and disassociation (such basic metaphors of cognition are discussed in Lakoff & Johnson 1980, and Barnden 1992, and given a computational treatment in Veale & Keane 1992a,b). The essential principle is this: when a speaker S utters a metaphor M to insult a sentient tenor T, such that B_HView(S, V) < 0, then H perceives S to *distance* himself conceptually from T; however, when a speaker S utters a metaphor M to praise a sentient tenor T, such that B_HView(S, V) > 0, then H perceives S to move himself conceptually *closer* to T.

If M is a metaphor uttered by S, describing T as V to H, and in doing so imparting the proposition set X to H (i.e., newly awakened bridges), then

For every newly awakened bridge T_X:V_X in X

- When M insults T, that is, B_HView(S, V) < 0

  When View(T, T_X) > View(T, V_X)
  i.e., Adopting T_X:V_X is detrimental to T
  infer: $\Delta B_H B_S B_T Z_{T_X}(V_X) \propto -Cf(M)$
  i.e., T distances himself from V_X in $B_H B_S$
  infer: $\Delta B_H B_S Z_{T_X}(V_X) \propto Cf(M)$
  i.e., S moves himself closer to V_X in $B_H B_S$

  When View(T, T_X) < View(T, V_X) > 0
  i.e., Adopting T_X:V_X is beneficial to T
  infer: $\Delta B_H B_S B_T Z_{T_X}(V_X) \propto Cf(M)$
  i.e., T moves himself closer to V_X in $B_H B_S$
  infer: $\Delta B_H B_S Z_{T_X}(V_X) \propto -Cf(M)$
  i.e., S distances himself from V_X in B_H

- When M praises T, that is, B_HView(S, V) > 0

  When View(T, T_X) > View(T, V_X)
  i.e., Adopting T_X:V_X is detrimental to T
  infer: $\Delta B_H B_S B_T Z_{T_X}(V_X) \propto -Cf(M)$
  i.e., T distances himself from V_X in $B_H B_S$
  infer: $\Delta B_H B_S Z_{T_X}(V_X) \propto -Cf(M)$
  i.e., S distances himself from V_X in B_H

  When View(T, T_X) < View(T, V_X)
  i.e., Adopting T_X:V_X is beneficial to T

infer: $\Delta B_H B_S B_T \; Z_{T_X}(V_X) \propto Cf(M)$

i.e., T moves himself closer to V_X in $B_H B_S$

infer: $\Delta B_H B_S \; Z_{T_X}(V_X) \propto Cf(M)$

i.e., S moves himself closer to V_X in B_H

Given these heuristics for belief update/revision in response to metaphor interpretation, the pejorative metaphor of Figure 1, "Surgeons are Butchers", can be analysed in terms of the speaker's desire to distance himself from a surgeon's beliefs. The system makes the natural assumption that a surgeon does not consider himself to be a butcher, in the interests of maintaining his self view, or esteem, and will therefore not lend any credibility to the comparison of surgery with slaughter, flesh with meat, or scalpels with cleavers. The speaker, however, in attacking surgeons, is perceived by the hearer to align himself fully with these beliefs.

7. Summary & Conclusions

Metaphor is an elegant and concise communicative form which is employed by a speaker as a means of conveying a state of affairs to a hearer; as such, it deserves to be analysed as a speech-act, with a particular illocutionary intent and perlocutionary effect. The state of affairs conveyed by metaphor is not always expressible in what is traditionally termed *literal language*, and thus the comprehension of metaphor is often a learning experience for the hearers, inasmuch as it requires them to reorganise their conceptual structures to accommodate the novel *analog-bindings* of the metaphor. Neither is this state of affairs always completely inherent in the propositional structure of the utterance; often much of the meaning conveyed by metaphor is pragmatic in nature, inasmuch as the metaphor provides the hearer with a glimpse into the belief-space of the speaker.

This paper has described an extension to the Sapper framework, a hybrid symbolic/connectionist model of metaphor previously outlined in Veale & Keane (1993), which incorporates elements of the belief framework of Wilks, Barnden & Wang (1991). This extended framework provides a suitable computational environment for analysing the illocutionary intent of the speaker, and perlocutionary effect upon the hearer, of a broad class of metaphors with an observable ameliorative/pejorative connotation. By concentrating on this class of metaphor, it is possible to speak of the overall *impression* imparted by a metaphor, avoiding commitment to particular propositional structures. This impression is easily quantified in terms of the activation dynamics of the memory network, and forms the basis for quantifying both the emotive and communicative force of a metaphor. These measures in turn underpin various belief ascription heuristics which allow the system to extract some of the pragmatic meaning conveyed by the metaphor.

This paper represents a first salvo in attacking the computational treatment of metaphor, not only as a propositional form, but as a communicative speech-act which both conveys speaker beliefs, and revises hearer beliefs. Future ground is yet to be made on this issue, however, in arriving at a model of metaphor which acknowledges the interplay between hearer and speaker belief-spaces as being essential to comprehension. Immediate extensions include investigations into a framework for quantifying *metaphoric aptness*, which we believe is a natural outgrowth of the current model.

References

Austin, J. L. (1962). *How to do things with words*, ed. J. Urmson, Oxford University Press.

Ballim, A., Y. Wilks & J. Barnden. (1991). Belief Ascription, Metaphor and Intensional Indentification. *Cognitive Science* 15(1), pp 133 - 171.

Barnden, J. A. (1992). Belief in Metaphor: Taking Commonsense Psychology Seriously, *Computational Intelligence*, 8(3).

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

Davidson, D. (1979). What Metaphors Mean, in The Philosophy of Language, ed. A. P. Martinich. Oxford University Press: London.

Finke, R. A., T. B. Ward & S. M. Smith (1992). *Creative Cognition: Theory, Research, and Applications. Cambridge*, MA: MIT Press.

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

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

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

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

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

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

Veale, T. & M. T. Keane. (1992a). Conceptual Scaffolding: Using metaphors to build knowledge structures, in *the Proc. of the 10th European Conference on Artificial Intelligence*, Chichester: John Wiley & Sons.

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

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

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

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

Wilks, Y., J. Barnden & J. Wang. (1991). Your Metaphor or Mine: Belief Ascription and Metaphor Interpretation, *in the Proc. of the 12th International Joint Conference on Artificial Intelligence*, IJCAI'91. San Mateo, CA: Morgan Kaufmann.