Affect Transfer by Metaphor for an Intelligent Conversational Agent

Alan Wallington, Rodrigo Agerri, John Barnden, Mark Lee, Tim Rumbell
School of Computer Science
University of Birmingham,
Birmingham B152TT
E-mail: A.M.Wallington@cs.bham.ac.uk

Abstract

We discuss an aspect of an affect-detection system used in edrama by intelligent conversational agents, namely affective interpretation of limited sorts of metaphorical utterance. Our system currently only deals with cases, which we found to be quite common in edrama, in which a person is compared to, or stated to be, something non-human such as an animal, object, artefact or supernatural being. Our approach permits a limited degree of variability and extension of these metaphors We discuss how these metaphorical utterances are recognized, how they are analysed and their affective content determined and in particular how the Electronic Lexical Database, WordNet, and the natural language glosses of the WordNet sysnsets can be used. We also discuss how this relatively shallow approach relates in important ways to the deeper ATT-Meta theory of metaphor interpretation and to approaches to affect and emotion in metaphor theory. We finish by illustrating the approach with a number of 'worked examples'.

1. Introduction

This paper discusses aspects of the extraction and processing of affective information such as emotions/ moods (e.g. embarrassment, hostility) and particularly evaluations (of goodness, importance, etc.) as conveyed by metaphor in free-form textual utterances. The background to this work is our experience in building upon an edrama system produced by one of our industrial partners, in which human users - school children, so far, in the testing and development stage of our work- improvise around certain themes by typing in utterances for the on-screen characters they play to utter (via speech bubbles). Drama by its nature involves emotional experience and this is furthered by the nature of the themes or scenarios we have used, namely 'school bullying' and a scenario involving a sufferer of a particularly embarrassing disease -Crohn's disease- discussing with friends and family whether or not to undergo an operation. User-testing (Zhang et al, 2006) shows that users have enjoyed using the system.

The need for the extraction and processing of affect arises because we have added to the edrama the option of having a bit-part character controlled by an Intelligent Conversational Agent (ICA). This ICA is capable of making largely contentless, but emotionally appropriate, interjections and responses in order to keep the conversation flowing, which it does by extracting affect from the human controlled characters' utterances. The same algorithms are also used for influencing the characters' gesturing (when a 3D animation mode produced by one of our industrial partners is used). Whilst other ICA research has concerned itself with the conveyance of affect (e.g. Picard, 2000), it appears that the conveyance of affect via metaphor has been largely ignored. Indeed, relatively little work has been done on any detailed computational processing of metaphor. Major exceptions include (Fass, 1997; Hobbs, 1990; Martin, 1990; Mason, 2004; Narayanan, 1999).

The background to the work on the conveyance of affect via metaphor comes from the authors' approach to, and partially implemented system (ATT-Meta) for, the processing and understanding of metaphor in general (Agerri et al. 2007; Barnden et al. 2004; Wallington et al. 2006). This is a more ambitious aim than the mere recognition of a metaphor or the classification of a metaphor into one of a number of different metaphor classes or conceptual metaphors (see Mason, 2004). The details of the implemented system need not concern us since they are not used in the control of the edrama ICA. However, aspects of the approach to metaphor are used. Thus, our metaphor approach and system emphasizes the open-endedness of metaphorical expressions, whereby conventional metaphors and fixed phraseology may be varied, extended and elaborated upon so as to convey further information and connotations not conveyed by the conventional metaphor. Although our ICA work uses WordNet for analysis of many of the affect-conveying metaphorical senses we find, we can analyse some phrasal variation in the words and deal with some senses that are not found.

Relatedly, our approach and system eschews large sets of correspondences between ontologically complex source and target domains in the manner of Lakoff and Johnson's (1980) 'Conceptual Metaphor Theory' e.g. ARGUMENT IS WAR, or ANGER IS HOT LIQUID UNDER PRESSURE (see Gibbs, 1992; Kövecses, 2000), with the meaning of a metaphorical utterance 'read off' from the source-target correspondences. Instead we assume very few, more abstract, specific source-target links between domains and account for much of the apparent systematic relatedness between source and target domains by noting that certain types of information, relations, attributes that can be inferred as holding of the situation described by a metaphorical utterance transfer in an invariant manner to

the target via a limited number of what we term View-Neutral Mapping Adjuncts or VNMAs. For example, we assume that if a causal link can be inferred as holding between entities in the source, then the causal link will hold by default in the target. Similarly, if something applies to a particular degree in the source, then its target equivalent will apply to the same degree in the source and likewise with such information as duration, temporal ordering, logical relations between entities, and others. Crucially for our edrama ICA, we assume that the emotional state that is either invoked by some aspect of the source, or that holds within the source will carry over to the target. We also assume that a value judgement concerning something in the source will also carry over by default to the target. For example consider a situation in which it is said of some foul mouthed character, 'Tom is a sewer'. This can be partially analysed in terms of Reddy's (1979/1993) well known 'conduit metaphor', in which information and utterances are viewed as if passing along a conduit from speaker to hearer, but crucially no source-target correspondence will be required for the specifics of 'sewer'. Instead, the negative value judgement about the nature of the material passing through a sewer should be transferred by the Value Judgement VNMA. A similar negative value judgement is conveyed by 'smelly attitude' or by the comment 'you buy your clothes at the rag market', two examples taken from transcripts the system automatically recorded during user-testing.

Emotional states and behaviour are often described metaphorically (Kövecses, 2000; Fussell & Moss, 1998), as in 'He was boiling inside' [feelings of anger] or 'He was hungry for her' [feelings of lust] and conceptual metaphors such as the above mentioned ANGER IS HOT LIQUID UNDER PRESSURE or LUST IS HUNGER proposed to account for this, but in an analysis of the transcripts from our user-testing the type of affect laden metaphor described in the previous paragraph was found to be a significant issue in edrama: at a conservative estimate, at least one in every 16 speech-turns has contained such metaphor (each turn is 100 characters, and rarely more than one sentence; 33000 words across all transcripts).

This paper will discuss how our system implements the transfer of affect in a very limited range of metaphors. However, it should be noted that the system underlying our edrama ICA does not detect affect solely or even primarily via metaphor. Quite apart from the recognition of specifically emotive and affective lexis, the system deals with letter and punctuation repetition for emphasis ("yeessss," "!!!!"), interjections and onomatopoeia (grrrrrrr) (see Zhang et al. 2006 for details). However, these may be viewed as manifestations of an abstract conceptual metaphor that views or conceptualises 'more of some thing or some quality' as 'an increase along one salient dimension'; typically height. This often gives us the Lakovian conceptual metaphor MORE IS UP, but gives word length when dealing with text. The degree of increase is conveyed by our degree VNMA.

Our system uses a blackboard architecture, in which hypotheses arising from the processing go onto a central blackboard. The production of the various hypotheses can then be influenced by hypotheses posted by other processes, etc. In particular, we envisage metaphor processing being refined by using such information (see Smith et al. 2007 for more details).

2. Affect via Metaphor in an ICA

Our system currently detects and analyses the transference of affect in the cases where a human is cast as a non-human of various sorts, as in the following cases:

1) Casting someone as an animal. This often transfers some affect -negative or positive- from the animal to the human. Interestingly, since our attitude towards young or baby forms, regardless of the animal concerned, are typically affectionate, affection is often transferred, even when the adult form is negative ('pig: piglet', 'dog: puppy' etc.). We deal with animal words that have a conventional metaphorical sense but also with those that do not, for it may still be possible to note a particular affective connotation, and even if not, one can plausibly infer that some affect or other is being expressed without knowing if positive or negative.

2) Relatedly, casting someone as a monster, mythical creature or supernatural being of some sort, using words such as 'monster' itself, 'dragon,' 'angel,' 'devil.'

3) Relatedly, casting someone as an artefact, substance or natural object, as in 'Tom is a [sewer; real diamond; rock].

We currently do not deal with the related case of casting someone metaphorically as a special type of human, using words such as 'baby,' 'freak,' 'girl' [to a boy], 'lunatic'.

In addition, size adjectives (cf. Sharoff 2006) often convey affect. Thus, 'a little X' can convey affective qualities of X such as an affectionate attitude towards X, even if the X is usually negative as in 'little devils' to describe mischievous children (compare with the baby forms above), but may sometimes convey unimportance and contemptibility as in 'you little rat'. Similarly, 'big X' can convey the importance of X ('big event') or intensity of X-ness ('big bully') -and X can itself be metaphorical as in 'big baby' when said of an adult.

3. Metaphor Processing

The approach is split into two parts: recognition of potential metaphors; analysis of recognised elements to determine affect. Note that in some cases, e.g. using 'pig' as a negative term for a person, the metaphor analysis requires only lexical look-up (e.g., in WordNet, 2006). But, not all animal words have a person sense and as noted above baby forms often change the affect as do size

adjectives. Such cases motivate the further processing.

3.1 The Recognition Component

The basis here is a list of words/phrases (www.cs.bham.ac.uk/jab/ATT-Meta/metaphoricity-signals.html) we term 'metaphoricity signals', that often have metaphors as collocates. They include specific syntactic structures as well as lexical strings. We currently focus on three syntactic structures, 'X is/are a Y', 'You Y' and 'like [a] Y' and on the lexical strings, 'a bit of a', 'such a' and 'look[s] like'. Note that a distinction is often made between similes and metaphors, making the third structure a simile. Our view is that (many) similes represent just a particular way of expressing an underlying metaphorical connection between X and Y and so shouldn't be treated differently from the other realisations. In the user-testing transcripts, we judged signals as actually involving metaphor in the following proportions of cases: *X is/are a Y* – 38% (18 out of 47); *you Y* – 61% (22 out of 36); *a bit of a / such a* – 40% (but tiny sample: 2 out of 5). Also: *looks like* and *like* – 81% (35 out of 43). (Of course, metaphor is often not signalled and can occur in any syntactic form and not just the forms here.)

In order to detect signals, the Grammatical Relations (GR) output from the RASP parser (Briscoe et al. 2006) is used. This output shows typed word-pair dependencies between the words in the utterance. For example, the following three GRs are output for a sentence such as 'You are a pig', so allowing an 'X is a Y' signal to be detected.

```
|ncsubj| |be+_vbr| |you_ppy| |_|
```

(i.e. the subject of 'are' is 'you')

```
|xcomp| | be+_vbr| |pig_nn1|
```

(i.e. the complement of 'are' is 'pig')

```
|det| |pig_nn1| |a_at1|
```

(i.e. the determiner of 'pig' is 'a')

Note that the tags 'vbr' and 'ppy' are specific to 'are' and 'you', so we also detect tags for: 'is'; for 'he', 'she' and 'it'; and for proper and common nouns, as well.

The output for the 'You Y' structure is typically as in the following example:

```
|ncmod| |you_ppy| |idiot_nn1|
```

(with Y = 'idiot') making it possible to find the structure from that one relation. But a common problem with RASP on 'You Y' is that its 'Part of Speech' (POS) tagger seems to favour tagging Y as a verb, if it can. For example, the word 'cow' in place of 'idiot' is tagged as a verb. In such a case, our system looks the word up in the list of tagged words that forms part of the RASP tagger. If the verb can be tagged as a noun, the tag is changed, and the metaphoricity signal is detected. Once a syntactic structure resulting from metaphoricity signals is detected, the word(s) in Y position are pulled out to be analysed.

This approach has the advantage that whether or not the noun in the Y position has adjectival modifiers the GR between the verb and Y is the same so the detection tolerates a large amount of variation, an important desiderata for metaphor. Any such modifiers are found in modifying relations and can be extracted for later analysis.

For additional confidence we detect the lexical strings 'a bit of a' and 'such a'. 'Such a' is found using GRs of the following type:

```
|det| |idiot_nn1| |an_at1|
```

(i.e. the determiner of 'idiot' is 'an'.)

```
|det| |idiot_nn1| |such_da|
```

(i.e. the determiner of 'idiot' is 'such')

Note that 'idiot', is detected as a 'Y' type metaphor, independently of 'such a', by the syntactic structure detection process: the 'X-is-a-Y' metaphoricity signal. The 'a bit of a' strings are found similarly, but cause the complication that the word 'bit' is tagged as a noun, so will be pulled out as a metaphor word by the syntactic detection processes, instead of the intended Y word. If the 'a bit of a' string is then found, we pull out the noun relating to the 'of' that relates to 'bit', in this type of GR output:

```
|iobj| |bit_nn1| |of_io|

|dobj| |of_io| |idiot_nn1|
```

In addition to 'X is a Y' and 'You Y', another metaphoricity signalling syntactic structure is 'like Y'. This is found using GR's of the following type:

```
|dobj| |like_ii| |pig_nn1|
```

'like Y' is always found in this form, with the noun in question in the dobj (direct object) relation to 'like', and with an nn1 tag. This is inserted into the list of present metaphoricity signals, and an additional flag is raised if it is found in an 'X looks like Y' structure. The 'looks like' structure can be uncovered by spotting this GR:

```
|iobj| |look_vv0| |like_ii|
```

Detection of the 'looks like' structure is similar to 'such a' in that it is in addition to the main metaphoricity signal detection, in this case not only adding confidence, but also potentially altering the meaning and analysis of the metaphor.

The result of the recognition element is threefold: (1) a list of signals; (2) the X and Y nouns from the syntactic signals; (3) a list of words modifying that noun.

3.2 The Analysis Component

The analysis element of the processing that we shall discuss here takes the X noun (if any) and Y noun and uses WordNet 2.0 (2006) to analyse them. First, we try to determine whether X refers to a person (the only case the system currently deals with), partly by using a specified list of proper names of characters in the drama and partly by WordNet processing (The system also proceeds similarly if X is 'you'). If so, then the Y and remaining elements are analysed using WordNet's taxonomy. This allows us to see if the Y noun in one of its senses is a hyponym of (or member of the class of) animals, supernatural beings, substances, artefacts or natural objects. If this is established, the system sees if another of the senses of the word is a hyponym of the person synset, as many metaphors are already given as senses in WordNet. If the given word contains different synsets or senses that are hyponyms of both animal etc. and person, then we search for evaluative content about the metaphor.

We have developed a method of automatically detecting the evaluation of a given metaphorical sense of a word. Intermediate synsets between the metaphorical sense of the given word and the person synsets contain glosses, which are descriptions of the semantic content of a synset. For example, the gloss of the synset of 'shark' that is a hyponym of 'person' is "a person who is ruthless and greedy and dishonest"; that of 'fox' is "a shifty deceptive person". We search the words and glosses from the intermediate synsets for words that indicate a particular affective evaluation. This is somewhat crude, since we do not parse the glosses, although a limited parser is currently being implemented. Consequently, both 'evil' and 'not evil' if found in a gloss will be taken to indicate a negative evaluation. (see Veale 2003 for related use of WordNet glosses).

Now there exist numerous lists and resources containing evaluative words. Indeed, SentiWordNet (Esuli & Sebastiani, 2006) is based on the glosses of the WordNet synsets and assigns three numerical scores describing how objective, positive, and negative the terms contained in the synset are. See also WordNet-Affect (Strapparava et al. 2004) However, in practice we found that very many of the animals etc. we wished to assign a positive or negative evaluation to were given a neutral score in SentiWordNet and so we created our own list. We decided that since we were searching though WordNet glosses, it would be most appropriate to create a list from WordNet itself. This we did in the following manner. WordNet contains a 'quality' synset which has 'attribute' links to four other synsets, 'good', 'bad', 'positive' and 'negative'. We are currently only looking for positive or negative affective evaluations, so this group of synsets provides a core set of affect indicating words to search for in the intermediate nodes. This set is expanded by following WordNet's 'see also' links to related words, to produce lists of positivity and negativity indicators. For example, 'bad' has 'see also' links to five synsets, including 'disobedient' and 'evil'; we then look up the 'see also' links in these five synsets and include these related words in the 'bad' list, and so on, through five iterations, producing a list of over 100 words indicating negativity.

With this list, we can search through the words and glosses from the intermediate nodes between the given metaphor synset (arising from the Y component in the sentence) and 'person', tallying the positivity and negativity indicating words found. We can then assign the affective evaluation of the metaphor, so having more negativity indicators than positivity indicators suggests that when the word is used in a metaphor it will be negative about the target. If the numbers of positivity and negativity indicators are equal, then the metaphor is labelled positive or negative, implying that it has an affective quality but we cannot establish what. This label is also used in those examples where an animal does not have a metaphorical sense in WordNet as a kind of person (for example, 'You elephant' or 'You toad').

It might be thought that the need for an additional person hypernym for Y is not necessary and that a search through the glosses of just the animal etc synsets in the hypernym tree Y would yield a relevant affective evaluation, at least in cases where there is no additional person sense. But this appears not to be the case. The glosses tend to be technical with few if any affective connotations. For example, 'toad' surprisingly does not have an alternative person sense in WordNet. The glosses of its 'amphibian, vertebrate and chordate hypernyms give technical information about habitat, breeding, skeletal structure, etc. but nothing affective. Worse still, false friends can be found. Thus, the word 'important' is used in many glosses in phrases like 'important place in the food chain' and this consequently causes some strange positive evaluations (for example of 'Cyclops' or 'water fleas').

We noted earlier that baby animal names can often be used to give a statement a more affectionate quality. Some baby animal names such as 'piglet' do not have a metaphorical sense in WordNet. In these cases, we check the word's gloss to see if it is a young animal and what kind of animal it is (The gloss for piglet, for example, is "a young pig").We then process the adult animal name to seek a metaphorical meaning but add the quality of affection to the result. A higher degree of confidence is attached to the quality of affection than is attached to the positive/negative result, if any, obtained from the adult name. Other baby animal names such as 'lamb' do have a metaphorical sense in WordNet independently of the adult animal, and are therefore evaluated as above. They are also tagged as potentially expressing affection, but with a lesser degree of confidence than that gained from the metaphorical processing of the word. However, the youth of an animal is not always encoded in a single word: e.g., 'cub' may be accompanied by specification of an animal type, as in 'wolf cub'. An extension to our processing would be required to handle this and also cases like

‘young wolf’ or ‘baby wolf’.

If any adjectival modifiers of the Y noun were recognized the analyser goes on to evaluate their contribution to the metaphor’s affect. If the analyser finds that ‘big’ is a modifying adjective of the noun it has analysed, the metaphor is marked as being more emphatic. If ‘little’ is found the following is done. If the metaphor has been tagged as negative and no degree of affection has been added (from a baby animal name, currently) then ‘little’ is taken to be expressing contempt. If the metaphor has been tagged as positive OR a degree of affection has been added then ‘little’ is taken to be expressing affection. These additional labels of affection and contempt are used to imply extra positivity and negativity respectively.

4. Examples of the Course of Processing

In this section we discuss three examples in detail and seven more with brief notes.

4.1 You piglet

1). The metaphor detector recognises the ‘You Y’ signal and puts the noun ‘piglet’ on the blackboard.
2). The metaphor analyser reads ‘piglet’ from the blackboard and detects that it is a hyponym of ‘animal’.
3). ‘Piglet’ is not encoded with a specific metaphorical meaning (‘person’ is not a hypernym). So the analyser retrieves the gloss from WordNet.
4). It finds ‘young’ in the gloss and retrieves all of the words that follow it. In this example the gloss is ‘a young pig’ so ‘pig’ is the only following word. If more than one word had followed, then the analysis process is repeated for each of the words following ‘young’ until an animal word is found
5). The words and glosses of the intermediate nodes between ‘pig’ and ‘person’ contain 0 positivity indicating words and 5 negativity indicating words, so the metaphor is labelled with negative polarity.
6). This example would result in the metaphor being labelled as an animal metaphor which is negative but affectionate with the affection label having a higher numerical confidence weighting than the negative label.

4.2 Lisa is an angel

1). The metaphor detector recognises the ‘X is a Y’ signal and puts the noun ‘angel’ on the blackboard. ‘Lisa’ is recognised as a person through a list of names provided with the individual scenarios in e-drama.
2). The metaphor analyser finds angel that it is a hyponym of ‘supernatural being’.
3). It finds that in another of its senses the word is a hyponym of ‘person’.
4). The words and glosses of the intermediate nodes between ‘angel’ and ‘person’ contain 8 positivity indicating words and 0 negativity indicating words, so the metaphor is labelled with positive polarity.
5). This example results in the metaphor being labelled as a positive supernatural being.

4.3 Mayid is a rock

1). The metaphor detector recognises the ‘X is a Y’ signal and puts the noun ‘rock’ on the blackboard. ‘Mayid’ is recognised as a person through a list of names provided with the individual scenarios in e-drama.
2). The metaphor analyser finds rock is a hyponym of ‘natural object’.
3) It finds that in another of its senses the word is a hyponym of ‘person’.
4). The words and glosses of the intermediate nodes between ‘rock’ and ‘person’ contain 4 positivity indicating words and 1 negativity indicating words, so the metaphor is labelled with positive polarity.
5). This example would result in the metaphor being labelled as a positive natural object.

4.4 Other Examples

1). ‘You cow’: this is processed as a negative animal metaphor. The synset of ‘cow’ that is a hyponym of ‘person’ has the gloss “a large unpleasant woman”. Interestingly, ‘large’ is included in the list of positivity indicators by the current compilation method, but the negativity of the metaphor is confirmed by analysis of the intermediate synsets between ‘cow’ and ‘person’, which are ‘unpleasant woman’, ‘unpleasant person’ and ‘unwelcome person’. These synsets, along with their glosses, contain six negativity and just one positivity indicator.
2). ‘You little rat’: this animal metaphor is determined as negative, having three senses that are hyponyms of ‘person’, containing three positivity indicators and five negativity indicators. ‘Little’ provides an added degree of contempt.
3). ‘You little piggy’: ‘piggy’ is recognized as a baby animal term and labelled as expressing affection. The evaluation of ‘pig’ adds a negative label, with no positivity indicators and three negativity indicators, and ‘little’ adds further affection since the metaphor already has this label from the baby animal recognition. This is therefore recognized as a negative metaphor but meant affectionately.
4). ‘You’re a lamb’: recognized as an animal metaphor and a young animal. It has an ‘affectionate’ label and is recognized as a positive metaphor, with its two senses that are hyponyms of ‘person’ contributing two positivity indicators and one negativity indicator. The negative word in this case is ‘evil’, coming from the gloss of one of the intermediate synsets, ‘innocent’: “a person who lacks knowledge of evil”. This example highlights a failing of using individual words as indicators: negations within sentences are not recognized.
5). ‘You are a monster’: one sense of monster in WordNet is a hyponym of animal. Therefore, this is recognized as an animal metaphor, but affect evaluation reveals three negativity and three positivity indicators, so it is analysed as ’positive or negative’. These indicators are found in two opposed senses of monster: ‘monster, fiend, ogre’: “a cruel wicked and inhuman person” (analysed as negative); and ‘giant, monster, colossus’: “someone that is

abnormally large and powerful" (analysed as positive, due to 'large' and 'powerful').
6). 'She's a total angel': a positive supernatural being metaphor, with eight positivity indicators and no negativity indicators from two senses that are hyponyms of 'person', but currently 'total' makes no contribution.
7). 'She is such a big fat cow': a negative animal metaphor made more intense by the presence of big. It has an extra level of confidence attached to its detection as two metaphoricity signals are present but currently 'fat' makes no contribution.

5. Conclusions and Further Work

The paper has discussed a relatively 'shallow' type of metaphor processing, although our use of robust parsing and complex processing of a thesaurus take it well beyond simple keyword approaches or bag-of-words approaches. Note that we do not wish simply to 'precompile' information about animal metaphor (etc.) by building a complete list of animals (etc.) in any particular version of WordNet (and also adding the effects of potential modifiers such as 'big' and 'little'), because we wish to allow the work to be extend to new versions of WordNet and to generalize as appropriate to thesauri other than WordNet, and because we wish to allow ultimately for more complex modification of the Y nouns, in particular by going beyond the adjectives 'big' and 'little'. We recognize that the current counting of positive and negative indicators picked up from glosses is an over-simple approach, and that the nature of the indicators should ideally be examined. This is a matter of both ongoing and future research. The processing capabilities described make particular but nonetheless valuable and wide-ranging contributions to affect-detection for ICAs. Although designed for an edrama system, the techniques plausibly have wider applicability. The development of the processing in a real-life application is also enriching our basic research on metaphor, such as the role of VNMAs.

6. Acknowledgements

This work has been supported by EPSRC grant EP/C538943/1 and grant RES-328-25-0009 from the ESRC under the ESRC/EPSRC/DTI "PACCIT" programme. We wish to thank our colleagues Catherine Smith and Sheila Glasbey.

7. References

Agerri, R., Barnden, J., Lee, M., Wallington, A. (2007). Metaphor, inference and domain independent mappings In *International Conference on Recent Advances in Natural Language Processing (RANLP 2007*, G Angelova, K Bontcheva, R Mitkov, N Nicolov, (Eds), Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 07), Borovets, Bulgaria, pp. 17-24.

Barnden, J., Glasbey, S., Lee, M., Wallington, A. (2004) Varieties and directions of interdomain influence in metaphor, *Metaphor and Symbol*, 19, pp.1-30.

Briscoe, E., Carroll, J. and Watson, R. (2006) The Second Release of the RASP System. *Proceedings of the COLING/ACL 2006 Interactive Presentation Sessions*. Sydney.

Esuli, A & Sebastiani, F. (2006). SentiWordNet: A Publicly Available Lexical Resource for Opinion Mining, In *Proceedings of the 5th Conference on Language Resources and Evaluation (LREC 2006)*, Genova, Italy.

Fass, D. (1997). *Processing Metaphor and Metonymy*. Greenwich, Connecticut: Ablex.

Fussell, S., & Moss, M. (1998). Figurative Language. In S.R. Fussell and R.J. Kreuz (Eds) *Emotional Communication. Social and Cognitive Approaches to Interpersonal Communication*. Hillside, NJ: Lawrence Erlbaum, pp. 113-142.

Gibbs, R. (1992) Categorization and metaphor understanding. *Psychological Review*. 99(3) 572-577

Hobbs, J. (1990). *Literature and Cognition*. Center for the Study of Language and Information, Stanford University: CSLI Lecture Notes, 21.

Martin, J. (1990). *A Computational Model of Metaphor Interpretation*. San Diego, CA: Academic Press.

Mason, Z. (2004). CorMet: A computational, corpus-based conventional metaphor extraction system. *Computational Linguistics*, 30(1).

Kövecses, Z. (2000). *Metaphor and Emotion: Language, Culture and Body in Human Feeling*. Cambridge: Cambridge University Press.

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago, IL: University of Chicago Press.

Narayanan, S. (1999). Moving right along: A computational model of metaphoric reasoning about events. In *Proceedings of the National Conference on Artificial Intigence*. AAAI Press.

Picard, R.W. (2000). *Affective Computing*. Cambridge MA: The MIT Press.

Reddy, M.J. (1979). The conduit metaphor: A case of frame conflict in our language about language. In A Ortony, (Ed) 2nd edn 1993 *Metaphor and Thought*. Cambridge: Cambridge University Press, 164-201.

Sharoff, S. (2006). How to Handle Lexical Semantics in SFL: a Corpus Study of Purposes for Using Size Adjectives. In S Hunston & G Thompson (Eds) *Systemic Linguistics and Corpus*. London: Equinox, pp 184-205.

Smith, C.J., Rumbell T.H., Barnden, J.A., Lee, M.G., Glasbey, S.R., & Wallington, A.M. (2007). Affect and Metaphor in an ICA: Further Developments. In Pelachaud, C.; Martin, J.-C.; André, E.; Chollet, G.; Karpouzis, K.; Pelé, D. (Eds.) *Intelligent Virtual Agents. 7th International Working Conference, IVA 2007*. Lecture Notes in Computer Science, Vol.4722. Heidelberg:,Springer-Verlag. pp 405-406

Strapparava, C. & Valitutti, V. (2004). WordNet-Affect: An Affective Extension of WordNet, In *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC 2004)* Lisbon, Portugal. 1083-1086.

Veale, T. (2003). Systematicity and the Lexicon in Creative Metaphor. In *Proceedings of the ACL Workshop on Figurative Language and the Lexicon, the 41st Annual Association for Computational Linguistics Conference* (ACL 2003), Sapporo, Japan.

Wallington, A, Barnden, J., Glasbey, S., Lee, M. (2006). Metaphorical reasoning with an economical set of mappings. *Delta* 22 (especial). pp 147-171

WordNet (2006). *WordNet, A Lexical Database for the English Language. Version 2.1* Cognitive Science Laboratory. Princeton University.

Zhang, L., Barnden, J.A., Hendley, R.J. & Wallington, A.M. (2006). Exploitation in Affect Detection in Improvisational E-drama. In J Gratch, M. Young, R. Aylett, D. Ballin, & P. Olivier (Eds) *Proceedings of the 6th International Conference on Intelligent Virtual Agents*. Lecture Notes in Computer Science, 4133, Springer, pp.68-79.