Don't worry about metaphor: affect extraction for conversational agents

Catherine Smith, Tim Rumbell, John Barnden, Bob Hendley, Mark Lee & Alan Wallington
School of Computer Science, University of Birmingham
Birmingham B15 2TT, UK
J.A.Barnden@cs.bham.ac.uk

Abstract

We demonstrate one aspect of an affect-extraction system for use in intelligent conversational agents. This aspect performs a degree of affective interpretation of some types of metaphorical utterance.

1 Introduction

Our demonstration is of one aspect of a system for extracting affective information from individual utterances, for use in text-based intelligent conversational agents (ICAs). Affect includes emotions/moods (such as embarrassment, hostility) and evaluations (of goodness, importance, etc.). Our own particular ICA [Zhang *et al.* 2006] is for use in an e-drama system, where human users behave as actors engaged in unscripted role-play. Actors type in utterances for the on-screen characters they control to utter (via speech bubbles). Our ICA is another actor, controlling a bit-part character. Through extracting affect from other characters' utterances it makes responses that can help keep the conversation flowing. The same algorithms are also used for influencing the characters' gesturing (when a 3D animation mode is used).

The system aspect demonstrated handles one important way in which affect is expressed in most discourse genres: namely metaphor. Only a relatively small amount of work has been done on computational processing of metaphorical meaning, for any purpose, let alone in ICA research. Major work apart from ours on metaphorical-meaning computation includes (Fass, 1997; Hobbs, 1990; Martin, 1990; Mason, 2004; Narayanan, 1999; Veale, 1998). The e-drama genre exhibits a variety of types of metaphor, with a significant degree of linguistic open-endedness. Also, note that our overarching research aim is to study metaphor as such, not just how it arises in e-drama. This increases our need for systematic, open-ended methods.

2 Metaphor and Affect

Conveying affect is one important role for metaphor, and metaphor is one important way of conveying affect. Emotional states and behavior often themselves described metaphorically (Kövecses, 2000; Fussell & Moss, 1998), as in 'He was boiling inside' [feelings of anger]. But another important phenomenon is describing something X using metaphorical source terms that are subject to that affect, as in 'My son's room [= X] is a *bomb site*' or '*smelly* attitude' (an e-drama transcript example). Such carry-over of affect in metaphor is well-recognized, e.g. in the political domain (Musolff, 2004). Our transcript analyses indicate that this type of affect-laden metaphor is a significant issue in e-drama: at a conservative estimate, in recent user studies in secondary schools at least one in every 16 speech-turns has contained such metaphor (each turn is $\leq$100 characters, and rarely more than one sentence; 33K words across all transcripts).

There are other specific, theoretically interesting metaphorical phenomena arising in e-drama that are important also for discourse in general, and plausibly could be handled reasonably successfully in an ICA using current techniques. Some are:
1) Casting someone as an animal. This often conveys affect, from insultingly negative to affectionately positive. Terms for young animals ('piglet', 'wolf cub', etc.) are often used affectionately, even

Proceedings of the ACL 2007 Demo and Poster Sessions, pages 37–40,
Prague, June 2007.

when the adult form is negative. Animal words can have a conventional metaphorical sense, often with specific affect, but in non-conventional cases a system may still be able to discern a particular affective connotation; and even if it cannot, it can still plausibly infer that *some* affect is expressed, of unknown polarity (positivity/negativity).

2) Rather similarly, casting someone as a monster or as a mythical or supernatural being, using words such as 'monster', 'dragon,' 'angel,' 'devil.'

3) Casting someone as a special type of human, using words such as 'baby' (to an adult), 'freak,' 'girl' (to a boy), 'lunatic.'

4) Metaphorical use of size adjectives (cf. Sharoff, 2006). Particularly, using 'a little X' to convey affective qualities of X such as unimportance and contemptibility, but sometimes affection towards X, and 'big X' to convey importance of X ('big event') or intensity of X-ness ('big bully')—and X can itself be metaphorical ('baby', 'ape').

Currently, our system partially addresses (1), (2) and (4).

3 Metaphor Recognition & Analysis

3.1 The Recognition Component

The basis here is a subset of a list of metaphoricity signals we have compiled [`http://www.cs.bham.ac.uk/~jab/ATT-Meta/metaphoricity-signals.html`], by modifying and expanding a list from Goatly (1997). The signals include specific syntactic structures, phraseological items and morphological elements. We currently focus on two special syntactic structures, *X is/are Y* and *You/you Y*, and some lexical strings such as '[looks] like', 'a bit of a' and 'such a'. The signals are merely uncertain, heuristic indicators. For instance, in the transcripts mentioned in section 2, we judged *X is/are Y* as actually indicating the presence of metaphor in 38% of cases (18 out of 47). Other success rates are: *you Y* – 61% (22 out of 36); *like* (including *looks like*)– 81% (35 out of 43).

In order to detect signals we use the Grammatical Relations (GR) output from the RASP robust parser [Briscoe *et al.*, 2006] This output shows typed word-pair dependencies between the words in the utterance. E.g., the GR output for 'You are a pig' is:

```
|ncsubj| |be+_vbr| |you_ppy| |_|
|xcomp| _ |be+_vbr| |pig_nn1|
|det| |pig_nn1| |a_at1|
```

For an utterance of the type *X is/are Y* the GRs will always give a subject relation (`ncsubj`) between X and the verb 'to be', as well as a complement relation (`xcomp`) between the verb and the noun Y. The structure is detected by finding these relations. As for *you Y*, Rasp also typically delivers an easily analysable structure, but unfortunately the POS tagger in Rasp seems to favour tagging Y as a verb—e.g., 'cow' in 'You cow'. In such a case, our system looks the word up in a 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 deemed detected. Once a signal is detected, the word(s) in relevant positions (e.g. the Y position) position are pulled out to be analysed. This approach has the advantage that whether or not the noun in, say, 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. Any such modifiers are found in modifying relations and are available for use in the Analysis Component.

3.2 The Analysis Component

We confine attention here to *X–is/are–Y* and *You–Y* cases. The analysis element of the processing takes the X noun (if any) and Y noun and uses WordNet 2.0 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. 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 animals or supernatural beings. 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 different senses of the given word are hyponyms of both animal and person, other categories in the tree between the noun and the person synset may provide information about the evaluative content of the metaphor. For example the word 'cow' in its metaphorical usage has the 'unpleasant person' synset as a lower hypernym, which heuristically suggests that, when the word is used in a metaphor, it will be negative about the target.

There is a further complication. 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 Word-

Net. In these cases, we check the word's gloss to see if it is a young animal and what kind of animal it is. 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 in the previous paragraph. They are also flagged 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 then goes on to evaluate their contribution to the metaphor's affect. If the analyser finds that 'big' is one of the modifying adjectives 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.

4 Examples of Course of Processing

'You piglet':

(1) Detector recognises the *you Y* signal with Y = 'piglet'.

(2) Analyser finds that 'piglet' is a hyponym of 'animal'.

(3) 'Piglet' does not have 'person' as a WordNet hypernym so analyser retrieves the WordNet gloss.

(4) It finds 'young' in the gloss ('a young pig') and retrieves all of the following words (just 'pig' – the analysis process is would otherwise be repeated for each of the words captured from the gloss), and finds that 'pig' by itself has negative metaphorical affect.

(5) The input is labelled as an animal metaphor which is negative but affectionate, with the affection having higher confidence than the negativity.

'Lisa is an angel':

(1) Detector recognises the *X is Y* signal with Y = 'angel', after checking that Lisa is a person.

(2) Analyser finds that 'angel' is a hyponym of 'supernatural being'.

(3) It finds that in another sense 'angel' is a hyponym of 'person'.

(4) It finds that the tree including the 'person' synset also passes through 'good person,' expressing positive affect.

(5) Conclusion: positive supernatural-being metaphor.

Results from Some Other Examples:

"You cow", "they're such sheep": negative metaphor.

"You little rat": contemptuous metaphor.

"You little piggy": affectionate metaphor with a negative base.

"You're a lamb": affectionate metaphor.

"You are a monster": negative metaphor.

"She is such a big fat cow": negative metaphor, intensified by 'big' (currently 'fat' is not dealt with).

5 Concluding Remarks

The demonstrated processing capabilities make particular but nevertheless valuable contributions to metaphor processing and affect-detection for ICAs, in e-drama at least. Further work is ongoing on the four specific metaphorical phenomena in section 3 as well as on other phenomena, such as the variation of conventional metaphorical phraseology by synonym substitution and addition of modifiers, and metaphorical descriptions of emotions themselves.

As many extensions are ongoing or envisaged, it is premature to engage in large-scale evaluation. Also, there are basic problems facing evaluation. The language in the e-drama genre is full of misspellings, "texting" abbreviations, acronyms, grammatical errors, etc., so that fully automated evaluation of the metaphorical processing by itself is difficult; and application of the system to manually cleaned-up utterances is still dependent on Rasp extracting structure appropriately. Also, our own ultimate concerns are theoretical, to do with the nature of metaphor understanding. We are interested in covering the qualitative range of possibilities and complications, with no strong constraint from their frequency in real discourse. Thus, statistical evaluation on corpora is not particularly relevant except for practical purposes.

However, some proto-evaluative comments that can be made about animal metaphors are as follows. The transcripts mentioned in section 2 (33K words total) contain metaphors with the following animal words: *rhino, bitch, dog, ape, cow, mole*, from 14 metaphorical utterances in all. Seven of the utterances are recognized by our system, and these involve *rhino, dog, ape, mole.* No WordNet-based metaphorical connotation is found for the *rhino* case. Negative affect is concluded for *bitch, dog* and *cow* cases, and affect of undetermined polarity is concluded for *ape* and *mole*.

The system is currently designed only to do relatively simple, specialized metaphorical processing. The system currently only deals with a small minority of our own list of metaphoricity signals (see section 3.1), and these signals are only present in a minority of cases of metaphor overall. It does not do either complex reasoning or analogical structure-matching as in our own ATT-Meta metaphor system (Barnden, 2006) or the cited approaches of Fass, Hobbs, Martin, Narayanan and Veale. However, we plan to eventually add simplified versions of ATT-Meta-style reasoning, and in particular to add the ATT-Meta *view-neutral mapping adjunct* feature to implement the default carry-over of affect (see section 2) and certain other information, as well as handling more signals.

Other work on metaphor has exploited WordNet (see, e.g., Veale, 2003, and panel on Figurative Language in WordNets and other Lexical Resources at GWC'04 (`http://www.fi.muni.cz/gwc2004/`. Such work uses WordNet in distinctly different ways from us and largely for different purposes. Our system is also distinctive in, for instance, interpreting the contribution of size adjectives.

Acknowledgments

The research has been aided by Sheila Glasbey and Li Zhang, and supported by ESRC/EPSRC/DTI Paccit LINK grant (ESRC RES-328-25-0009) and EPSRC grant EP/C538943/1.

References

John Barnden. 2006. Artificial Intelligence, Figurative Language and Cognitive Linguistics. In G. Kristiansen *et al.* (Eds), *Cognitive Linguistics: Current Applications and Future Perspectives,* 431–459. Berlin: Mouton de Gruyter.

Ted Briscoe, John Carroll and Rebecca Watson. 2006. The Second Release of the RASP System. In *Procs. COLING/ACL 2006 Interactive Presentation Sessions*. Sydney, Australia.

Dan Fass. 1997. *Processing Metaphor and Metonymy.* Greenwich, Connecticut: Ablex.

Susan Fussell & Mallie Moss. 1998. Figurative Language in Emotional Communication. *Social and Cognitive Approaches to Interpersonal Communication*. Lawrence Erlbaum.

Andrew Goatly. 1997. *The Language of Metaphors*. Routledge, London.

Jerry Hobbs. 1990. *Literature and Cognition.* CSLI Lecture Notes, 21, Stanford University, 1990.

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

James Martin. 1990. *A Computational Model of Metaphor Interpretation.* Academic Press.

Zachary Mason. 2004. CorMet: A computational, corpus-based conventional metaphor extraction system. *Computational Linguistics*, 30(1), 23–44.

Andreas Musolff. 2004. *Metaphor and political discourse: Analogical reasoning in debates about Europe.* Palgrave Macmillan.

Srini Narayanan. 1999. Moving right along: A computational model of metaphoric reasoning about events. *Procs. National Conference on Art. Int.*, 121–128.

Serge Sharoff. 2006. How to Handle Lexical Semantics in SFL: A Corpus Study of Purposes for Using Size Adjectives. *System and Corpus: Exploring Connections*. Equinox, London.

Tony Veale. 1998. 'Just in Time' analogical mapping, an iterative-deepening approach to structure-mapping. In *Procs. 13th European Conference on Art. Intell.*

Tony Veale. 2003. Dynamic Type Creation in Metaphor Interpretation and Analogical Reasoning: A Case-Study with WordNet. In *Procs. International Conference on Conceptual Structures* (Dresden).

Li Zhang, John Barnden, Bob Hendley & Alan Wallington. 2006. Exploitation in Affect Detection in Improvisational E-Drama. In Procs. 6th Int. Conference on Intelligent Virtual Agents: *Lecture Notes in Computer Science, 4133*, 68–79. Springer.