

Affect in Metaphor: Developments with WordNet

Tim Rumbell, John Barnden, Mark Lee and Alan Wallington¹

Abstract. We discuss an aspect of an affect-detection system used in e-drama by intelligent conversational agents, namely affective interpretation of limited sorts of metaphorical utterance. We discuss how these metaphorical utterances are recognized and how they are analysed and their affective content determined.

1 INTRODUCTION

We present 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 [19] 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). Our ICA is an addition to an e-drama system produced by an industrial collaborator, Hi8us Midlands Limited, and the 3D animation facility was produced by another industrial collaborator, BT. See [4] for more.

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 ([6], [9], [11], [12], [14], [17]). 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. This paper updates [15]. Since that system-demonstration paper important new developments have taken place in the metaphor processing.

2 METAPHOR AND AFFECT

Conveying affect is one important role for metaphor, and metaphor is one important way of conveying affect. Emotional states and behaviour are often themselves described metaphorically ([10]; [7]), 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 ([13]). We recently conducted a user study of the system

(at four secondary schools in the Birmingham area) and (automatically) recorded the different users' (actors') "speeches". An analysis of the resulting transcripts indicate that this type of affect-laden metaphor is a significant issue in e-drama: 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; 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 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) sheerly from the fact of using an animal metaphor.
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'. These again can have strong (if context-sensitive) affective connotations.
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'). Metaphorical use of size adjectives is often combined with phenomena 1-3. In particular, "little X" where X is an animal can lead to a baby-animal interpretation as one possibility.

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

3 METAPHOR RECOGNITION AND ANALYSIS

The approach is split into two parts: recognition of potential metaphors; analysis of recognised elements to determine affect.

3.1 The Recognition Component

The basis here is a subset of a list of metaphoricity signals we have compiled [1], by modifying and expanding a list from [8]. The signals include specific syntactic structures, phraseological items and

¹ University of Birmingham, UK, email: a.m.wallington@cs.bham.ac.uk

morphological elements. We currently focus on two special syntactic structures, *X is/are Y* (in which X could be the pronoun ‘you’ and *You Y*, and some lexical strings such as ‘[looks] like’, ‘a bit of a’ and ‘such a’ (these lexical strings can be interpreted as indicative of similes, but we treat similes and metaphors in the same way). 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 per cent of cases (18 out of 47). Other success rates are: *you Y* - 61 per cent (22 out of 36); *like* (including *looks like*) - 81 per cent (35 out of 43).

In order to detect signals we use the Grammatical Relations (GR) output from the RASP robust parser [3]. This output shows typed wordpair 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’. Here the robustness of a parser like RASP causes problems: the main verb sense of ‘cow’ (overawe: “subdue, restrain, or overcome by affecting with a feeling of awe”) is transitive and so would not normally be relevant to ‘You cow’, but it is desirable for a robust parser to allow deficient grammar. 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 metaphoricality signal is deemed detected.

Once a signal is detected, the word(s) in relevant positions (e.g. the Y 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

3.2.1 Core Processing

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 (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 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 the given word contains within its senses different senses that are hyponyms of both animal and person, then we search for evaluative content about the metaphor.

Previously our analysis of evaluative content of a metaphor revolved around finding specific indicative synsets in the hypernym tree of the person sense of the given metaphor word. For example,

cow in its metaphorical sense has the ‘unpleasant person’ synset as a lower hypernym than ‘person’, which we took as an indicator of negativity). But now, in an important new development, instead of relying on the presence of one of a small set of intermediate nodes for affective evaluation, we have developed a method of automatically detecting the orientation of a given metaphorical 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 search is based on another feature of WordNet, as follows.

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 related to ‘bad’, and therefore indicating negativity. We 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 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 labeled 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’). See the comment at the end of case 1 in section 2.

3.2.2 Young Animals and Size Adjectives

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 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 by means of the Core Processing in section 3.2.1. They are also tagged as potentially expressing affection but with a lesser degree of confidence than that gained from the core 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. These additional labels of affection and contempt are used to imply extra positivity and negativity respectively.

4 EXAMPLES OF COURSE OF PROCESSING

In this section we discuss two examples in detail and seven more with brief notes. The examples are mainly from [15] but we have updated several to take account of the gloss-based processing mentioned in section 3.2.1. The first two, the detailed examples, outline the flow of processing and highlight the key analytical decisions made.

4.1 'You piglet'

1. The metaphor detector recognises the *You Y* signal and tags the noun 'piglet' as the Y word.
2. The metaphor analyser reads 'piglet' from as Y 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 is following, 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 labeled 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 tags the noun 'angel' as the metaphor word. 'Lisa' is recognised as a person through a list of names provided with the individual scenarios in e-drama.
2. The metaphor analyser finds 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 labeled with positive polarity.
5. This example would result in the metaphor being labeled as a supernatural being metaphor that is positive.

4.3 Other examples

The following are further examples to show some of the ways in which particular types of utterance are analysed.

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 indicators, against just the 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 labeled 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 currently 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 the indicators '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 metaphoricality signals are present but currently 'fat' makes no contribution.

5 FUTURE WORK

Work is ongoing on the four specific metaphorical phenomena listed in section 2 as well as on other phenomena, such as the variation of conventional metaphorical phraseology by synonym substitution and addition of modifying words and phrases, and interpretation of metaphorical descriptions of emotions. We are also looking to broaden metaphor detection, such that, in the case *X is/are Y*, if a hypernym of X is an 'artifact' and of Y is a 'living thing' (or vice-versa) then a metaphor is implied.

Observe 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’). This is because we wish to allow the work to be extended 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.

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. In future work we wish to integrate the processing we have described with the deep semantic/pragmatic reasoning-based approach in our ATT-Meta project [2]. Note also that the carry over of affect in animal (etc.) metaphor as treated above is a special case of a much more general carry-over phenomenon that is central to the ATT-Meta approach (cf. its “view-neutral mapping adjuncts” feature).

Our reason for not reporting full evaluations at this stage is the amount of extensions ongoing or envisioned. We have extra work to do on the system and as such it is premature to engage in a large scale evaluation.

6 RELATED WORK

WordNet glosses have been used elsewhere to extract additional information about metaphor. Veale [18] describes an approach to qualia extraction from WordNet glosses, by attempting to extract relational structure inherent from them. Similarities between this approach to glosses and our own can be found in that glosses are used to determine compatibility between concepts, finding structures not explicitly encoded in WordNet already. Veale also highlights the use of finding relevant analogues of particular concepts to find relations with other concepts, similar to our method of relating words from glosses to the concepts of positivity and negativity. Of course, Veale uses these techniques to understand similarities in relational structure in metaphors and we use them only to determine an affective component. On the other hand, our processing is for extracting information from sentences, whereas Veale’s is not directly applied to this.

Other systems of affectively labeling WordNet synsets have been developed. SentiWordNet [5] and WordNet-Affect [16] both annotate synsets with affective labels, and SentiWordNet uses the information available in glosses to do so. Transferring our work to work in these systems could be useful; our processing could potentially fill gaps in these systems. Our processing could still be useful were this not the case, because of our interest, mentioned above, in generalizing the work to non-WordNet resources.

7 CONCLUSIONS

The processing capabilities described make particular but nonetheless valuable and wide-ranging contributions to affect-detection for ICAs. Although designed for an e-drama system, the techniques plausibly have wider applicability. That is, it is both the case that animal, supernatural-being and big/little metaphors appear in many other genres, and that the techniques we have developed for such metaphor can plausibly be generalized to work for a variety of other

types of metaphor. The development of the processing in a real-life application is also enriching our basic research on metaphor.

ACKNOWLEDGEMENTS

The research was supported in part by an ESRC/EPSRC/DTI Pacific LINK grant (ESRC RES-328-25-0009) and by EPSRC grant EP/C538943/1. We are indebted to our industrial partners (Hi8us Midlands Ltd, Maverick TV Ltd, and BT), to our colleagues Li (Jane) Zhang and Catherine Smith for intensive system-development work, and to other colleagues on the e-drama project (Rodrigo Agerri, Sheila Glasbey, Bob Hendley and William Edmondson).

REFERENCES

- [1] <http://www.cs.bham.ac.uk/jab/ATT-Meta/metaphoricity-signals.html>.
- [2] R. Agerri, J. Barnden, M. Lee, and A. Wallington, *Metaphor, Inference and Domain-Independent Mappings*, Proceedings of the Int. Conf. Recent Advances in Natural Language Processing (RANLP’07), Borovets, Bulgaria, 2007.
- [3] E. Briscoe, J. Carroll, and R. Watson, *The Second Release of the RASP System*, Proceedings of the COLING/ACL 2006 Interactive Presentation Sessions, Sydney, 2006.
- [4] K. Dhaliwal, M. Gillies, J. O’Connor, A. Oldroyd, D. Robertson, and L. Zhang, *Drama: Facilitating Online Role-play Using Emotionally Expressive Characters*, 195–202, Artificial and Ambient Intelligence: Proceedings of the AISB Annual Convention, 2007.
- [5] A. Esuli and F. Sebastiani, *SentiWordNet: A Publicly Available Lexical Resource for Opinion Mining*, Proceedings of the 5th Conference on Language Resources and Evaluation (LREC 2006), Genova, Italy, 2006.
- [6] D. Fass, *Processing Metaphor and Metonymy*, Ablex, Greenwich, Connecticut, 1997.
- [7] S. Fussell and M. Moss, *Figurative Language in Emotional Communication*, 113–135, Social and Cognitive Approaches to Interpersonal Communication, Lawrence Erlbaum Associates, 1998.
- [8] A. Goatly, *The Language of Metaphors*, Routledge, London, 1997.
- [9] J. Hobbs, *Literature and Cognition*, 21, CSLI Lecture Notes, Centre for the Study of Language and Information, Stanford University, 1990.
- [10] Z. Kövecses, *Metaphor and Emotion: Language, Culture and Body in Human Feeling*, Cambridge University Press, 2000.
- [11] J. Martin, *A Computational Model of Metaphor Interpretation*, Academic Press, San Diego, CA, 1990.
- [12] Z. Mason, ‘A computational, corpus-based conventional metaphor extraction system’, *Computational Linguistics*, **30(1)**, 23–44, (2004).
- [13] A. Musolf, *Metaphor and political discourse: Analogical reasoning in debates about Europe*, Palgrave Macmillan, Basingstoke, UK, 2004.
- [14] S. Narayanan, *Moving Right Along: A Computational Model of Metaphoric Reasoning About Events*, 121–18, Proceedings of the National Conference on Artificial Intelligence, AAAI Press, 1999.
- [15] C. Smith, T. Rumbell, J. Barnden, B. Hendley, M. Lee, and A. Wallington, *Don’t Worry About Metaphor: Affect Extraction for Conversational Agents*, Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics (ACL-2007), 2007.
- [16] C. Strapparava and V. Valitutti, *WordNet-Affect: An Affective Extension of WordNet*, 1083–1086, Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC 2004), Lisbon, Portugal, 2004.
- [17] T. Veale, ‘Just in Time’ Analogical Mapping, *An Iterative-Deepening Approach to Structure-Mapping*, Proceedings of the 13th European Conference on Artificial Intelligence (ECAI’98), John Wiley, 1998.
- [18] T. Veale, *Systematicity and the Lexicon in Creative Metaphor*, Proceedings of the ACL Workshop on Figurative Language and the Lexicon, the 2003 Applied Computational Linguistics Conference (ACL 2003), Sapporo, Japan, 2003.
- [19] L. Zhang, J. Barnden, B. Hendley, and A. Wallington, *Exploitation in Affect Detection in Improvisational E-Drama*, 68–79, Lecture Notes in Computer Science, Springer, Amsterdam, 2006. Also as Report No. 17/92, Computing Laboratory, University of Kent at Canterbury, U.K.