

Affect Transfer by Metaphor for an Intelligent Conversational Agent

Alan Wallington and Rodrigo Agerri and John Barnden and Mark Lee and Tim Rumbell

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 synssets 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’.

Alan Wallington

School of Computer Science, University of Birmingham, Birmingham B152TT, e-mail: Alan.Wallington@dunelm.org.uk, A.M.Wallington@cs.bham.ac.uk

Rodrigo Agerri

School of Computer Science, University of Birmingham, Birmingham B152TT, e-mail: R.Agerri@cs.bham.ac.uk

John Barnden

School of Computer Science, University of Birmingham, Birmingham B152TT, e-mail: J.A.Barnden@cs.bham.ac.uk

Mark Lee

School of Computer Science, University of Birmingham, Birmingham B152TT, e-mail: M.G.Lee@cs.bham.ac.uk

Tim Rumbell

School of Computer Science, University of Birmingham, Birmingham B152TT, e-mail: T.H.Rumbell@cs.bham.ac.uk

1 Introduction

In this paper we discuss aspects of the extraction and processing of affective information such as emotions/ moods (e.g. embarrassment, hostility) and evaluations (of goodness, importance, etc.) when conveyed by metaphor in free-form textual utterances. This work follows from our experience in building upon an edrama system produced by an industrial partner of ours. Using this system, human users –school children, in the testing and development stage of our work– improvised around certain themes by typing in utterances for on-screen characters they play to utter (via speech bubbles).

Now, drama by its very nature involves emotional experience, and this is particularly the case in our use of the edrama system because of the nature of the themes or scenarios we 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. The need for the extraction and processing of this emotion and affect arises because we added to the edrama the option of having a bit-part character controlled by an Intelligent Conversational Agent (ICA). This ICA can intervene and make largely contentless, interjections and responses in order to keep the conversation flowing. And, the interjections are emotionally appropriate because the underlying system can extract affect from the human controlled characters’ utterances¹. Whilst other ICA research has concerned itself with the conveyance of affect [15], 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 [5, 8, 12, 13, 14].

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 [1, 2, 22]. 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 [13]). The details of the implemented system need not concern us since they are not used in the control of the edrama ICA. However, particular important aspects of the ATT-Meta approach, if not the ATT-Meta system, are used in this much more applied, edrama, system. 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.

¹ User-testing [25] shows that users have enjoyed using the system.

² 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.

Along with, and relating to, our stress on the potential open-endedness of many metaphorical expressions, our approach and system eschews large sets of correspondences between ontologically complex source and target domains in the manner of Lakoff and Johnson's [11] 'Conceptual Metaphor Theory' e.g. ARGUMENT IS WAR, or ANGER IS HOT LIQUID UNDER PRESSURE (see [7, 10]), with the meaning of a metaphorical utterance 'read off' from the source-target correspondences. Instead we assume very few, more abstract, source-target links between specific 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 apparently being described under the source or literal interpretation of a metaphorical utterance transfer in an essentially invariant manner from the source, where they were inferred, to the target. This is done via a limited number of what we term View-Neutral Mapping Adjuncts or VNMA, with a particular VNMA for each type of information that transfers in this invariant manner. 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, transferred to the target via the CAUSATION VNMA. Similarly, if something can be viewed as applying to a particular degree in the source, e.g. slowly, then by virtue of the DEGRE VNMA, its target equivalent will be understood as applying to the same degree and likewise with such information as duration, temporal ordering, logical relations between entities, and others. Crucially for our edrama ICA, we note two particular VNMA. Any emotional state that is invoked either by some aspect of the source, or that holds within the source, will carry over to the target by virtue of the EMOTIONAL STATE VNMA. We also assume that a value judgement concerning something in the source will also carry over by default to the target via the VALUE JUDGEMENT VNMA. 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 [16] well known 'conduit metaphor', in which information and utterances are viewed as if passing along a conduit from speaker to hearer, but what about the negative value judgement denoted by 'sewer'? Crucially no source-target correspondence specific to the conduit metaphor will be required for this information. Instead, the negative value judgement about the nature of the material passing through a sewer should be transferred by the Value Judgement VNMA to become a negative value judgement about the words being used or spoken. 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.

Now, within metaphor theory, it has long been noted that emotional states and behaviour are often described metaphorically [10, 6], 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. However, the conveyance of emotion via VNMA is of a different type to such specific metaphors, and in an analysis of the transcripts from our user-testing, the type of affect laden metaphor that motivates the emotional state and value-judgement VNMA was

found to be a significant issue in edrama: at a conservative estimate, at least one in every 16 speech-turns has contained such a metaphor (each turn is 100 characters, and rarely more than one sentence; 33000 words across all transcripts).

We shall now 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 (grrrrrr) (see [25] for details). Although, note that these too 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.

Finally, note that 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 [18] for more details).

2 Affect via Metaphor in an ICA

As noted, our system currently only applies to a limited type of affect transfer; it detects and analyses the transference of affect in cases where a human is metaphorically 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.) and this seems quite productive. 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, or to a woman], ‘lunatic’.

We noted the modification of the affect conveyed by animal metaphors when the young form is used. In addition, size adjectives [17] 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 [24]). 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 (see www.cs.bham.ac.uk/jab/ATT-Meta/metaphoricity-signals.html, [23]) we term ‘metaphoricity signals’, that often have metaphors as collocates. They include specific syntactic structures as well as lexical strings. Of course, metaphor is often not signalled and can occur in any syntactic form and not just the forms in these lists. Here, we 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).

In order to detect signals, the Grammatical Relations (GR) output from the RASP parser [3] is used. This output shows typed word-pair dependencies between the words in the utterance. Using the grammatical relations from RASP 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. Consequently, 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 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.

For additional confidence with this syntactic, ‘X-is-a-Y’, metaphoricity signal, we also detect the lexical strings ‘a bit of a’ and ‘such a’, as in ‘he is such an idiot’. Note ‘idiot’ would have been detected as a ‘Y’ type metaphor, independently of ‘such a’, by the syntactic structure detection process.

‘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’)

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|
```

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. However, we found a problem with using RASP to detect these ‘You Y’ structures. It was frequently the case that RASP’s ‘Part of Speech’ (POS) tagger will prefer to tag ‘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.

A third metaphoricity signalling syntactic structure we use 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 [24] to analyse them. First, we try to determine whether X refers to a person (the only case the system currently deals with). One possibility is if the X is a pronoun such as ‘you’. However, we also use a specified list of proper names of characters in the drama. If X is a person, then the Y and remaining elements are analysed using WordNet’s taxonomy.

The system first checks whether the Y noun in one of its synsets (or senses) is a hyponym of (or member of the class of) animals, supernatural beings, substances, artefacts or natural objects. If this is established, then the following procedure applies.

- (1) The system tries to determine whether another of the senses of the word is a hyponym of the person synset, as many such ‘person as animal etc.’ metaphors are already given as separate senses in WordNet. If the given word contains different synsets or senses that are hyponyms of both animal etc. and person, we search for evaluative content about the metaphor and assign the result to Y
- (2) If there are no additional person synsets, or if no positive or negative evaluation can be determined, then the system will look down the various hyponym chains of the word to find all instances that also have a sense that can be determined to be a type of person (i.e. are hyponyms of the person synset). The evaluative content of these is then sought and a count made of the negative versus positive evaluations. Currently, if the ratio of one to the other is 3 or greater, then the Y is given a positive or negative evaluation accordingly.

- (3) If this search too results in failure, then the hypernym chain of the word is searched up to the ‘animal, artefact, etc, level, regardless of whether they have a person sense, and these are evaluated as either positive or negative.

So how do we determine positive or negative evaluations for the different senses? 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 (or the animal, artefact, etc synsets if all else fails) contain glosses, which are free-text descriptions of the semantic content of a synset. Take, for example, the word ‘shark’. One of its synsets is a hyponym of ‘person’. The gloss of this particular synset states that it is “a person who is ruthless and greedy and dishonest”. Similarly, one synset of ‘fox’ has a gloss as “a shifty deceptive person”. We search the words and glosses from the intermediate synsets for words, such as ‘ruthless’ or ‘shifty’ that indicate a particular affective evaluation. (see [21] for related use of WordNet glosses).

But, how is it known that ‘ruthless’, for example, has a negative evaluation? Now there exist numerous lists and resources containing evaluative words. Indeed, SentiWordNet [4] 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 [20]. 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 through 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. If this procedure is followed through four iterations, it produces a list of 37 ‘Negative’ words, 18 ‘Positive’ words, 68 ‘Bad’ words and 77 ‘Good’ words.

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.

However, if this were all that is done, then an incorrect evaluation would be assigned when the glosses include negation with ‘not’. For example, the gloss for ‘persona non grata’ is ‘a person who for some reason is not wanted or welcome’ and ‘wanted’ and ‘welcome’ are on the list of positive words. What we need instead are the terms ‘unwanted’ or ‘unwelcome’. To deal with such cases, we employ a procedure called ‘NotChecker’. This searches the gloss for structures of the type ‘is not X’ or ‘is not X or Y’. Once it has found a structure like this, it will look in WordNet for an antonym of X (and Y, if necessary), antonyms are listed in WordNet, and the X and Y are replaced by their antonyms when the search is made for positive

or negative evaluations. For example, in the 'persona non grata' case, NotChecker will detect the 'not wanted or welcome' in the gloss and the WordNet entries for 'wanted' and 'welcome' are examined for antonyms. Note to guard against multiple antonyms with different senses, the antonyms are themselves checked for antonyms to see if the original word is listed and it is this antonym that is used. An antonym of 'wanted' is 'unwanted' and so 'unwanted' is used in the evaluation search.

Given the words from the glosses, including those added via NotChecker, and the lists of affective words derived by following 'see also' links from 'good' and 'bad', we can then assign the affective evaluation of the metaphor. If there are more negativity indicators than positivity indicators, it 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'). In these cases, case (2) of the procedure described above will be tried.

It might be thought that the need for an additional person hypernym for Y is not necessary and that case (3) of the procedure on its own will suffice, i.e. a search through the glosses of just the animal etc synsets in the hypernym tree for Y would yield a relevant affective evaluation. 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'). It is for this reason that this search is carried out only as a last resort.

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. Finally, note that 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 the NotChecker procedure.

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.
8. ‘He’s a reptile’: reptile has as animal sense and no person sense, but has 4 negative hyponyms and no positive hyponyms, so it is evaluated as negative through the hyponym method (i.e. 2).

5 Results

It must be borne in mind when evaluating the system, exactly what the system described here is being used for. This is to detect and determine the affect if any in the utterances of human controlled avatars in an edrama setting, when a non-human quality is being attributed to a human. The system was not designed to return an accurate affectual score to any possible animal, artefact, etc that might be given to the system, or to populate WordNet or similar with such scores. Thus the fact that the system evaluates, perhaps counter-intuitively, water fleas as conveying positive, rather than neutral, affect is of little concern unless water fleas were likely to be used with the intention of conveying some affect in utterances such as ‘You are a water flea’.

With this in mind we can report on some evaluations of the system. There are a total of 27053 entries in WordNet for the five categories of animal, artefact, natural object, substance and spiritual being. To create a gold standard against which the performance of the system could be compared, one of the authors went through this list choosing items which were judged to have a strong positive or negative evaluation. This list was then circulated for comment or revision and a total of 63 positive and 141 negative entities decided upon.

We noted in Subsection 3.2 that we created a list of positive and negative evaluative terms by iteratively following WordNet ‘see also’ links from the terms ‘good’ and ‘bad’. Clearly, the more iterations made the larger the list, but the more likely it is that some of the words on the list would not express a positive or negative evaluation. After some experimentation we settled on a list created by four iterations.

We also noted in Subsection 3.2 that we would give an evaluation via the hyponym method if the ratio of positives to negatives or vice-versa was 3 to 1. Using these two parameters in the system, a total of 3832 items were given a positive evaluation and 1128, a negative evaluation. Of the 3832 positive evaluations, 41 appeared on the gold standard list, i.e. 65.1% of the gold standard positive items were correctly tagged. Of the 1128 items given a negative evaluation, 59 appeared on the gold standard list, i.e. 42% were correctly tagged. Of course, conversely of the 3832 positive evaluations, only 1% (41) appeared on the gold standard positive list and of the 1128 negative evaluations, only 5.2% were also on the gold standard negative list. But as noted, for our purposes these latter statistics are not very important.

We noted earlier that we had found SentiWordNet [4], which assigns numerical scores indicating how objective, positive or negative WordNet senses are, to be inadequate for our needs, and we can give some results to support this claim, by also evaluating it against our gold standard. Of the 63 positive senses, SentiWordNet only gave a higher positive than negative score to 14 of these, against 41 for our system. Indeed, it treated 4 of them as negative: ‘cherub’; ‘rock’; ‘Christ’; ‘sweet’. Items receiving no evaluation included such seemingly positive entities as ‘angel’ ‘lion’ and ‘diamond’. Turning to negatives, both our system and SentiWordNet assigned a negative evaluation to 59 items. However, there were 8 items SentiWordNet marked as positive, including ‘crap’, ‘Judas’ and ‘sheep.’ Among the items receiving no positive or negative evaluation were ‘toad’, ‘vermin’ and ‘punk’.

6 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 ontologies and 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 VNMA. In particular our use of hypernym trees in WordNet and potentially in other ontologies, chimes with contemporary views of metaphor theory that emphasise category mem-

bership, such as the work of [9] or recent work in Relevance Theory and metaphor [19]³. However, our search for specific types of information, as with our VNMA approach, that might be found in the hypernym trees is not found in such approaches.

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.

References

1. Agerri, R. and Barnden, J. and Lee, M. and Wallington, A. (2007). Metaphor, inference and domain independent mappings (pp. 17–24). In: *Angelova, G. and Bontcheva, K. and Mitkov, R. and Nicolov, N. (eds.), Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 07)*, Borovets, Bulgaria.
2. Barnden, J. and Glasbey, S. and Lee, M. and Wallington, A. (2004). Varieties and directions of interdomain influence in metaphor. *Metaphor and Symbol*, vol. 19, pp. 1–30.
3. Briscoe, E. and Carroll, J. and Watson, R. (2006). The Second Release of the RASP System (pp. 77–80). In: *Proceedings of the COLING/ACL 2006 Interactive Presentation Sessions*, Sydney.
4. Esuli, A. and Sebastiani, F. (2006). SentiWordNet: A Publicly Available Lexical Resource for Opinion Mining (pp. 417–422). In: *Proceedings of the 5th Conference on Language Resources and Evaluation (LREC 2006)*, Genova, Italy.
5. Fass, D. (1997). *Processing Metaphor and Metonymy*. Greenwich, Connecticut: Ablex.
6. Fussell, S. and Moss, M. (1998). Figurative Language (pp. 113–142). In: Fussell, S. R. and Kreuz, R. J. (eds.), *Emotional Communication. Social and Cognitive Approaches to Interpersonal Communication*. Hillsdale, NJ: Lawrence Erlbaum.
7. Gibbs, R. (1992) Categorization and metaphor understanding. *Psychological Review*, vol. 99(3), pp. 572–577.
8. Hobbs, J. (1990). *Literature and Cognition*. Center for the Study of Language and Information, Stanford University: CSLI Lecture Notes, 21.
9. Glucksberg, S. and Keysar, B. (1993). How metaphors work (pp. 401–424). In: Ortony, A. (Ed.), *Metaphor and thought* (2nd ed.). Cambridge, Cambridge University Press.
10. Kövecses, Z. (2000). *Metaphor and Emotion: Language, Culture and Body in Human Feeling*. Cambridge: Cambridge University Press.
11. Lakoff, G. and Johnson, M. (1980). *Metaphors we live by*. Chicago, IL: University of Chicago Press.
12. Martin, J. (1990). *A Computational Model of Metaphor Interpretation*. San Diego, CA: Academic Press.
13. Mason, Z. (2004). CorMet: A computational, corpus-based conventional metaphor extraction system. *Computational Linguistics*, vol. 30(1), pp. 23–44.
14. Narayanan, S. (1999). Moving right along: A computational model of metaphoric reasoning about events (pp. 121–127). In: *Proceedings of the National Conference on Artificial Intelligence*. AAAI Press.
15. Picard, R. W. (2000). *Affective Computing*. Cambridge MA: The MIT Press.
16. Reddy, M. J. (1979). The conduit metaphor: A case of frame conflict in our language about language (pp. 164–201). In: Ortony, A. (Ed.), *Metaphor and Thought* (2nd ed. 1993). Cambridge: Cambridge University Press.

³ Although we did not pursue this in our work, exploring meronym and holonym links in WordNet might be of use in interpreting limited types of metonymy.

17. Sharoff, S. (2006). How to Handle Lexical Semantics in SFL: a Corpus Study of Purposes for Using Size Adjectives (pp. 184–205). In: Hunston, S. and Thompson, G. (eds.), *Systemic Linguistics and Corpus*. London: Equinox.
18. Smith, C. J. and Rumbell T. H. and Barnden, J. A. and Lee, M. G. and Glasbey, S. R. and Wallington, A. M. (2007). Affect and Metaphor in an ICA: Further Developments (pp. 405–406). In: Pelachaud, C. and Martin, J.-C. and André, E. and Chollet, G. and Karpouzis, K. and Pelé, D. (eds.), *Intelligent Virtual Agents. 7th International Working Conference, IVA 2007*. Lecture Notes in Computer Science, Vol.4722. Heidelberg: Springer-Verlag.
19. Sperber, D. and Wilson, D. 2008. A deflationary account of metaphor (pp. 84–109). In: Gibbs, R. (Ed.), *Handbook of Metaphor and Thought*. Cambridge, Cambridge University Press.
20. Strapparava, C. and Valitutti, V. (2004). WordNet-Affect: An Affective Extension of WordNet (pp. 1083–1086). In: *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC 2004)*, Lisbon, Portugal.
21. Veale, T. (2003). Systematicity and the Lexicon in Creative Metaphor (pp. 28–35). In: Wallington, A. M. (Ed.), *Proceedings of the ACL Workshop on Figurative Language and the Lexicon, the 41st Annual Association for Computational Linguistics Conference (ACL 2003)*, Sapporo, Japan.
22. Wallington, A. and Barnden, J. and Glasbey, S. and Lee, M. (2006). Metaphorical reasoning with an economical set of mappings. *Delta*, vol. 22 (especial), pp. 147–171.
23. Wallington, A. M. and Barnden, J. A. and Barnden, M. A. and Ferguson, F. J. and Glasbey, S. R. (2003). *Metaphoricity Signals: A Corpus-Based Investigation*. CSRP-03-5, School of Computer Science, Birmingham University.
24. WordNet (2006). *WordNet, A Lexical Database for the English Language*. Version 2.1 Cognitive Science Laboratory. Princeton University.
25. Zhang, L. and Barnden, J. A. and Hendley, R. J. and Wallington, A. M. (2006). Exploitation in Affect Detection in Improvisational E-drama (pp. 68–79). In: Gratch, J. and Young, M. and Aylett, R. and Ballin, D. and Olivier, P. (eds.), *Proceedings of the 6th International Conference on Intelligent Virtual Agents. Lecture Notes in Computer Science*, 4133, Springer.