

Exploitation in Affect Detection in Improvisational E-drama

Li Zhang, John A. Barnden, Robert J. Hendley, and Alan M. Wallington

School of Computer Science, University of Birmingham, Birmingham, B15 2TT, UK
l.zhang@cs.bham.ac.uk

Abstract. We report progress on adding affect-detection to a program for virtual dramatic improvisation, monitored by a human director. To aid the director, we have partially implemented emotion detection within users' text input. The affect-detection module has been used to help develop an automated virtual actor. The work involves basic research into how affect is conveyed through metaphor and contributes to the conference themes such as building improvisational intelligent virtual agents for interactive narrative environments.

Keywords: E-drama, affect detection, intelligent virtual actor and metaphor.

1 Introduction

Improvised drama and role-play is widely used in education, training, counselling and conflict resolution and researchers are exploring frameworks for e-drama, in which virtual characters (avatars) on computer displays interact under the partial or total control of human users (e.g.[1]). Our research stemmed from an existing system, *edrama*, created by Hi8us Midlands Ltd and used variously in schools for e.g. creative writing. Many students welcome the anonymity of *edrama*. One main aspect of our project is the addition of types of intelligent automation.

In the *edrama* system, up to five virtual characters are controlled on a virtual stage by human users ("actors"), with characters' (textual) "speeches" typed by the actors operating the characters. A director is also involved in a session. A graphical interface on each actor's and director's terminal shows the stage and characters. Speeches are shown as text bubbles. Actors choose their characters clothes and bodily appearance.

Currently, cartoon figures against backdrops of real-life photographic images are used. However, we are bringing in animated gesturing avatars and 3D computer-generated settings using technology from an industrial partners, BT. Actors and the human director work through software clients connecting with the server. Clients communicate using XML stream messages via the server, which is usually remote from the terminals, which may themselves be remote from each other. Terminal-server communication is over the Internet using standard browsers.

The actors are given a loose scenario around which to improvise, but are at liberty to be creative. One scenario we have used is school-bullying where a schoolgirl Lisa is being bullied by her classmate Mayid. There are also roles for two friends and a teacher. Within these parameters, actors must improvise interesting interchanges.

The human director has a number of roles. S/he must constantly monitor the unfolding drama and the actors' interactions, or lack of them, in order to check whether they are keeping to the general spirit of the scenario. If this is not happening, the director may then intervene. For example, a director may intervene when the emotions expressed or discussed by characters are not as expected (or are not leading consistently in a new interesting direction). The director may also feel the need to intervene if one character is not getting involved, or is dominating the improvisation.

Intervention can take a number of forms. The director can send messages to actors. However, another important means of directorial intervention is for the director to introduce and control a 'bit-part' character. This character will not have a major role in the drama, but might, for example, try to interact with a character who is not participating much in the drama or who is being ignored by the other characters. Alternatively, it might make comments intended to 'stir up' the emotions of those involved, or, by intervening, diffuse any inappropriate exchange developing.

Clearly, all this places a heavy burden on the director. In particular, playing the role of the bit-part character and interacting with other characters whilst keeping interventions limited so as to maintain the main improvisatory drama amongst the actors, makes it difficult to fully monitor the behaviour of all the other actors and send appropriate messages to them should they stray off topic or exhibit inappropriate emotions. The difficulty is particularly acute if the directors are novices, such as teachers trying to use e-drama in their lessons.

One major research aim is accordingly to automate some directorial functions, either to take some of the burden away from a human director, or to provide a fully automated (though necessarily very restricted) director. With a fully-automated director, even if highly restricted in what it can do, little or no human supervision might be required for at least minimally adequate improvisations, and *edrama* could, for example, be added to websites about certain topics allowing visitors to engage in on-line role-play germane to the topic.

However, our main current work is on assisting a human director by providing fully-automated control of a bit-part character (though we are also working on automating limited types of director-to-actor message-sending to allow the human director to concentrate on the more difficult aspects of the task). For this reason, we have created an automated actor, EMMA (emotion, metaphor and affect), which operates a bit-part character (e.g. an acquaintance of the main character), and is under the control of an affect-detection module. The module tries to identify affect in other characters' speeches, allowing the EMMA character to make responses that will hopefully stimulate the improvisation. Within affect we include: basic and complex *emotions* such as anger and embarrassment respectively; *meta-emotions* such as desiring to overcome anxiety; *moods* such as hostility; and *value judgements* (of goodness, importance, etc.). Although merely detecting affect is limited compared to extracting the full meaning of characters' utterances, we have found that in many cases this is sufficient for the purposes of stimulating the improvisation.

Also, even limited types of affect detection are useful. EMMA may not detect all types of affect under all ways it can be expressed or implied, nor do it with a high degree of reliability, but the spirit of the project is to see how far we can get with practical processing techniques, while at the same time investigating theoretically the

nature of, and potential computational ways of dealing with, forms of affective expression that are too difficult to handle in a usable implemented system.

Much research has been done on creating affective virtual characters in interactive systems. Picard's work [2] makes great contributions to building affective virtual characters. Also, emotion theories, particularly that of Ortony, Clore and Collins [3] (OCC), have been used widely therein. Prendinger and Ishizuka [4] used the OCC model to reason about emotions and to produce believable emotional expressions. Wiltschko's *eDrama Front Desk* [5] is an online emotional natural language dialogue simulator with a virtual reception interface for pedagogical purposes. Mehdi et al. [6] combined the widely accepted five-factor model of personality [7], mood and OCC in generating emotional behaviour for a fireman training application. Gratch and Marsella [8] presented an integrated model of appraisal and coping, in order to reason about emotions and to provide emotional responses, facial expressions and potential social intelligence for virtual agents. Egges, Kshirsagar and Magnenat-Thalmann [9] have provided virtual characters with conversational emotional responsiveness. Elliott, Rickel and Lester [10] demonstrated tutoring systems that reason about users' emotions. There is much other work in a similar vein.

There has been only a limited amount of work directly comparable to our own, especially given our concentration on improvisation and open-ended language. However, *Facade* [11] included shallow natural language processing for characters' open-ended utterances, but the detection of major emotions, rudeness and value judgements is not mentioned. Zhe and Boucouvalas [12] demonstrated an emotion extraction module embedded in an Internet chatting environment (see also [13]). It uses a part-of-speech tagger and a syntactic chunker to detect the emotional words and to analyse emotion intensity for the first person (e.g. 'I' or 'we'). Unfortunately the emotion detection focuses only on emotional adjectives, and does not address deep issues such as figurative expression of emotion. Also, the concentration purely on first-person emotions is narrow. We might also mention work on general linguistic clues that could be used in practice for affect detection (e.g. [14]).

Our work is distinctive in several respects. Our interest is not just in (a) the first-person, positive expression of affect case: the affective states or attitudes that a virtual character X implies that it itself has (or had or will have, etc.), but also in (b) affect that the character X implies it lacks, (c) affect that X implies that other characters have or lack, and (d) questions, commands, injunctions, etc. concerning affect. We aim also for the software to cope partially with the important case of communication of affect via metaphor [15, 16], and to push forward the theoretical study of such language, as part of our research on metaphor generally (see, e.g., [17]).

Our project does not involve using or developing deep, scientific models of how emotional states, etc., function in cognition. Instead, the deep questions investigated are on linguistic matters such as the metaphorical expression of affect. In studying how ordinary people understand and talk about affect in ordinary life, what is of prime importance is their *common-sense* views of how affect works, irrespective of scientific accuracy of those views. Metaphor is strongly involved in such views.

It should also be appreciated that this paper does not address the emotional, etc. states of the *actors* (or director, or any audience). Our focus is on the affect that the actors make their characters express or mention. While an actor may work him/herself up into, or be put into, a state similar to or affected by those in his/her own characters'

speeches or those of other characters, such interesting effects, which go to the heart of the dramatic experience, are beyond the scope of this paper, and so is the possibility of using information one might be able to get about actors' own affective states as a hint about the affective states of their characters or vice-versa.

2 Our Current Affect Detection

Various characterizations of emotion are used in emotion theories. The OCC model uses emotion labels (anger, etc.) and intensity, while Watson and Tellegen [18] use positivity and negativity of affect as the major dimensions. Currently, we use an evaluation dimension (negative-positive), affect labels, and intensity. Affect labels plus intensity are used when strong text clues signalling affect are detected, while the evaluation dimension plus intensity is used when only weak text clues are detected. Moreover, our analysis is based on the transcripts of previous e-drama sessions. Since even a person's interpretations of affect can be very unreliable, our approach combines various weak relevant affect indicators into a stronger and more reliable source of information for affect detection. Now we summarize our affect detection based on multiple streams of information.

2.1 Pre-processing Modules

The language in the speeches created in e-drama sessions severely challenges existing language-analysis tools if accurate semantic information is sought even in the limited domain of restricted affect-detection. The language includes misspellings, ungrammaticality, abbreviations (often as in text messaging), slang, use of upper case and special punctuation (such as repeated exclamation marks) for affective emphasis, repetition of letters or words for emphasis, and open-ended interjective and onomatopoeic elements such as "hm" and "grrrr". In the examples we have studied, which so far involve teenage children improvising around topics such as school bullying, the genre is similar to Internet chat.

To deal with the misspellings, abbreviations, letter repetitions, interjections and onomatopoeia, several types of pre-processing occur before actual detection of affect.

A lookup table deals with abbreviations (e.g. 'im (I am)' and 'c u (see you)'). Most abbreviations used in Internet chat rooms and textese can be so handled. We also deal with abbreviations such as numbers embedded within words, e.g., "18r" for later using the lookup table. We handle the ambiguity of, say, "2" (to, too, two) in textese (e.g. "I'm 2 hungry 2 walk"), by using two simple rules that consider the POS tags of immediately surrounding words. In evaluations using examples in previous transcripts we have obtained 85.7% accuracy, which is adequate currently.

Letter repetition comes in two flavours: repetition added to ordinary words (e.g. 'yesss', 'seee'); and repetition added to interjections or onomatopoeic elements (e.g. 'grrrrr', 'agggghhh'). The iconic use of written word length here (corresponding roughly to imagined sound length) normally implies strong affective states in the characters' input. Usefully, adding letters does not greatly change the pronunciation. We have a small dictionary containing base forms of interjections (e.g. 'grr') and

some ordinary words that often have letters repeated in e-drama. Then the Metaphone spelling-correction algorithm [20], which is based on pronunciation, works with the dictionary to locate the base forms of words with letter repetitions. We also aim to develop a detector of onomatopoeic elements that does not rely on particular base forms. We must stress that added letter-repetition is not simply eliminated, but the fact of its occurrence is recorded for the purposes of affect-detection.

Finally, the Levenshtein distance algorithm [21] with a contemporary English dictionary deals with spelling mistakes in users' input. Having described the necessary pre-processing, we turn to the core detection of affect in users' input.

2.2 Affect Detection by Pattern Matching

In an initial stage, we based affect detection purely on textual pattern-matching rules that looked for simple grammatical patterns or templates partially involving lists of specific alternative words. This continues to be a core aspect of our system but we have now added robust parsing using Rasp [19] and some semantic analysis.

In textual pattern-matching, particular keywords, phrases and fragmented sentences are found, but also certain partial sentence structures are extracted. This procedure possesses the robustness and flexibility to accept many ungrammatical fragmented sentences and to deal with the varied positions of sought-after phraseology. However, it lacks other types of generality and can be fooled when the phrases are suitably embedded as subcomponents of other grammatical structures. For example, if the input is "I doubt she's really angry", rules looking for anger in a simple way will fail to provide the expected results. Below we indicate our path beyond these limitations.

The transcripts analysed to inspire our initial knowledge base and pattern-matching rules had independently been produced earlier from Hi8us' *edrama* improvisations. The actors were school children aged from 8 to 12. We have also worked on another, distinctly different scenario - Crohn's disease, based on a TV programme about this disease by Maverick Television Ltd. (another of our industrial partners). One interesting feature in this scenario is meta-emotion (emotion about emotion) and cognition about emotion, because of the need for people to cope with emotions about their illnesses. The rule sets created for one scenario have a useful degree of applicability to other scenarios, though there will be a few changes in the related knowledge database according to EMMA's different roles in specific scenarios.

A rule-based Java framework called Jess [22] is used to implement the pattern-matching rules in EMMA. When Mayid says "Lisa, you Pizza Face! You smell", EMMA detects that he is insulting Lisa. Patterns such as 'you smell' have been used for rule implementation. The rules conjecture the character's emotions, evaluation dimension (negative or positive), politeness (rude or polite) and what response EMMA should make. Here is one simple pseudo-code example rule.

```
(defrule example_rule
  ?fact <-(any string containing 'get out')
=>
  (obtain emotion and response from knowledge database)
```

When a character says “Lisa, get out of here” and EMMA responds, this example rule will be fired. EMMA infers the affective quality from the utterance (*angry and rude* in this case) and obtains the appropriate response from the knowledge database.

Multiple exclamation marks and the capitalisation of whole words are frequently employed to express emphasis in e-drama sessions. If these are detected in a character’s utterance, then the emotion intensity is deemed to be comparatively high (and emotion is suggested even in the absence of other indicators).

A reasonably good indicator that an inner state is being described is the use of ‘I’ (see also Craggs and Wood [14]), especially in combination with the present or future tense. In the school-bullying scenario, when ‘I’ is followed by a future-tense verb the affective state ‘threatening’ is normally being expressed; and the utterance is usually the shortened version of an implied conditional, e.g., “I’ll scream [if you stay here].” When ‘I’ is followed by a present-tense verb, other emotional states tend to be expressed, e.g. “I want my mum” (fear) and “I hate you” (dislike). Further analysis of first-person, present-tense cases is described in section 2.4.

2.3 Processing of Imperatives

A useful signal of strong emotions is the imperative mood, especially when used without softeners such as ‘please’ or ‘would you’. We deal with some common imperative phrases such as “shut up” and “mind your own business” directly. They often indicate strong negative emotions. But the phenomenon is more general.

Detecting imperatives accurately is by itself an example of the non-trivial problems we face. To go beyond the limitations of the text matching we have done, we have also used syntactic outputs from the Rasp parser and semantic information in the form of the semantic profiles for the 1,000 most frequently used English words [23] to deal with certain types of imperatives. This helps us to deal with some of the difficulties.

The Rasp parser recognises some types of imperatives directly. Unfortunately, the grammar of the 2002 version that we have used does not deal properly with certain imperatives (John Carroll, p.c). This means that examples like “you shut up”, “Matt don’t be so blunt”, “please leave me alone” and “don’t you call me a dog”, are not recognized as imperatives, but as declaratives or interrogatives. Hence, extra analysis is needed to detect imperatives, by additional processing applied to the possibly-incorrect syntactic trees Rasp produces. This includes consideration of the nature of the sentence subject, the form of the verb used and whether negation is present. We mention one case of special interest, involving semantic and pragmatic processing.

When a sentence involves a subject and a verb for which there is no difference at all between the base form and the past tense form, then an imperative/declarative ambiguity arises (e.g. “Lisa hit me”). An important special case of this ambiguity is when the object of the verb is ‘me’. In order to solve the ambiguity, we have adopted the evaluation value of the verb from Heise’s compilation of semantic differential profiles [23]. In these profiles, Heise listed values of evaluation, activation, potency, distance from neutrality, etc. for the 1,000 most frequently used English words. In the evaluation dimension, positive values imply goodness. Because normally people tend to use ‘a negative verb + me’ to complain about an unfair fact, if the evaluation value of the verb is negative, then the sentence is probably not an imperative but a statement

sentence (e.g. “Mayid hurt me”). Otherwise, other factors implying imperative are checked in this sentence, such as exclamation marks and capitalizations. If these factors occur, then the input is probably an imperative. Otherwise, the conversation logs are checked to see if there is any question sentence directed toward this speaker recently. If there is, then the input is conjectured to be declarative.

2.4 Using WordNet for a First Person Case

As we mentioned earlier, the first-person with a present-tense verb tends to express an affective state in the speaker. We have used the Rasp parser to detect such a sentence. We are exploiting WordNet [24] synsets to broaden the average of the existing pattern-matching rules to allow rough synonyms of the verbs in the rules to lead to a similar affect, and the analysis of synsets is refined by using Heise’s [23] semantic profiles. For example, if the user’s input is “I enjoy the movie very much”, we use WordNet to obtain the synonyms of the verb ‘enjoy’ (possibly from different synsets). The set of synonyms is refined by using semantic profiles from Heise’s dictionary and we obtain rough synonyms ‘love’ and ‘like’. Then we use ‘love’ to replace the verb ‘enjoy’, and send the newly built sentence “I love the movie very much” to the pattern-matching rules in order to obtain the speaker’s affective state and EMMA’s response. If we cannot successfully obtain such information, we will build another input sentence using the other synonym ‘like’ and send the sentence “I like the movie very much” to the pattern-matching rules. In general, using Wordnet provides us with the benefit of making our affect detection approach more generalized.

2.5 Responding Regimes

EMMA normally responds, on average, to every Nth speech by another character, where N is a changeable parameter (currently usually set to 3). However, it also responds when EMMA’s character’s name is mentioned, and makes no response if nothing useful is detected in the utterance. The one-in-N average is achieved by sampling a random variable every time another character says something, although in the future, we plan to have N dynamically adjustable according to how confident EMMA is about what it discerns in the utterance at hand. The random response is chosen from several stored response candidates that are suitable for the affective quality discerned in the utterance being responding to. For example,

<Another character> Don’t mess with me. Don’t forget it!

<EMMA character> Hey, tone down your language a bit.

In addition, EMMA may take part of the user’s input string as a sub-component of its response. Because the pre-processing module, reported in section 2.1, replaces the abbreviations and misspellings with ‘translations’ of the original, this helps obscure the fact that part of EMMA’s response is only a reflection. For example,

<Another character> I’m ur (*your*) best m8 (*mate*).

<EMMA character> Why are you my best mate?

On the other hand, we purposely put abbreviations and other phenomena used in Internet chat rooms in responses in order to make EMMA ‘behave’ like a school kid.

For example,

<Another character> You get lost!

<EMMA character> Hey, coool it folks. (notice the lengthening of 'cool')

<Another character> He didn't do it.

<EMMA character> No1 believes you, m8. (where 'no1' stands for 'no one')

The reader should bear in mind that EMMA's responses are not aimed at engaging with the detailed meaning of the utterance, but simply to stimulate the improvisation in a way that is somewhat sensitive to affect being expressed or mentioned. Furthermore, in order to make the EMMA character's responses push the improvisation forward, the character will not only ask scenario related questions to the main characters, but also introduce new topics closely related to the scenario in the improvisation. In a recent user-testing debrief session, secondary school students mentioned that the human bit-part character did not stay in character and said pointless things, while in another session one student, who played a main character, believed that the EMMA character was the only one that stuck to scenario related topics. The directors reported that, even when a main character was silent and the director did not intervene very much, the EMMA character led the improvisation on the right track by raising new topics other characters were concerned about.

3 Affect via Metaphor

Metaphor is relevant to affect detection for two reasons. It is often used to convey affect and it partly underlies folk theories of affect and emotion. For example, folk theories often talk about, and appear to conceive of, anger as if it were a heated fluid possibly exerting a strong pressure on its containing body. This motivates a wide range of metaphorical expressions both conventional such as "he was boiling with anger and about to blow his top" and more creative variants such as "the temperature in the office was getting higher and this had nothing to do with where the thermostat was set" (modified, slightly from a Google™ search). Passion, or lack of, is also often described in terms of heat and the latter example could in certain contexts be so used.

So far, examples of actors reflecting or commenting on the nature of their or others emotions have been infrequent in the e-drama transcripts, although we might expect to find more examples as more students participate in the Crohn's disease scenario. However, such metaphorically motivated folk models often directly motivate the terminology used to convey affect, as in utterances such as "you leave me cold", which conveys lack of interest or disdain. This use of metaphor to motivate folk models of emotions and, as a consequence, certain forms of direct expression of emotion has been extensively studied in linguistics, [15], [16].

Less recognised is the fact that metaphor frequently conveys emotion more indirectly. Here the metaphor does not describe some aspect of an emotional state, but something else. Crucially, however, it also conveys a negative or positive value judgement which is carried over to what is being described and this attitude hints at the emotion. For example to say of someone's room that "it is a cess-pit" allows the negative evaluation of 'cess-pit' to be transferred to 'the room' and we might assume an emotion of disgust. In our transcripts we find examples such as "smelly attitude"

and “you buy your clothes at the rag market” (which we take to be not literally true). Animal insults such as “you pig” frequently take this form, although many are now highly conventionalised. Our analysis of e-drama transcripts shows that this type of metaphor is much more common than the direct use.

It should be apparent that even though conventional metaphorical phraseology may well be listed in specialised lexicons, approaches to metaphor and affect which rely upon a form of lexical look-up to determine the meaning of utterances are likely to miss both the creative variants and extensions of standard metaphors and also the quite general carrying over of affectual evaluations from the literal meaning of an utterance to the intended metaphorical meaning.

At the time of writing (May/June 2006) EMMA incorporates little in the way of metaphor handling. However, certain aspects will be incorporated shortly, since they involve extensions of existing capabilities. Our intended approach is partly to look for stock metaphorical phraseology and straightforward variants of it, which is the most common form of metaphor in most forms of discourse, including edrama. We also plan to employ a simple version of the more open-ended, reasoning-based techniques described in the ATT-Meta project on metaphor processing ([17] [25] [26]).

As a first step, it should be noted that insults and swear words are often metaphorical. We are currently investigating specialised insult dictionaries and the machine-readable version of the OALD, which indicates slang.

Calling someone an animal of any sort often conveys affect, but it can be insulting or affectionate. The young of an animal is often used affectionately, and the same is true of diminutive (e.g., ‘piglet’) and nursery forms (e.g., ‘moo cow’), even when the adult form is usually used as an insult. Thus calling someone ‘a cat’ or ‘catty’ differs from describing them as kittenish. Likewise, “you pup” differs from “you dog”. We are constructing a dictionary of specific animals used in slang and as insults, but for animals not listed we can use WordNet and electronic dictionaries to determine whether or not it is the young or mature form of the animal that is being used.

We have already noted that in metaphor the affect associated with a source term by default carries across to the target. EMMA already consults Heise’s compilation of semantic differential profiles for the evaluation value of the verb. We will extend the determination of the evaluation value to all parts of speech.

Having the means to determine the emotion conveyed by a metaphor is most useful when metaphor can be spotted reliably. There are a number of means of doing this for some metaphors. Thus, idioms are often metaphorical [27] and we can use an existing idiom dictionary, adding to it as necessary. Unfortunately, as is often noted, idioms often show a degree of variation, either by adding modifiers, e.g., ‘shut your **big fat** mouth’ or by using synonyms of standard lexis, e.g., ‘**constructing** castles in the air’ instead of ‘building castles in the air’. This variability will pose a challenge if one is looking for fixed expressions. However, if the idiom dictionary is treated as providing base forms, with for example the nouns being treated as the heads of noun-phrases, then the Rasp parser can be used to determine the noun phrase and the modifiers of the head noun, and likewise with verbs, verb-phrases, etc. Indeed, this approach can be extended beyond highly fixed expressions to other cases of metaphor, since as [28] has noted metaphors tend to display a much greater degree of fixedness compared to non-metaphors, whilst not being as fixed as what are conventionally called idioms.

There are other ways of detecting metaphors. Metaphoricity signals ([29] [30]) such as: *so to speak*, *sort of*, *almost* may signal metaphor use. Semantic restriction violations as in “my car **drinks** petrol,” often indicate, metaphor ([31] [32] [33]), although not all metaphors violate semantic restrictions. To determine whether semantic restrictions are being violated, domain information from ontologies/thesauri such as WordNet could be used and/or statistical techniques as used by Mason ([33]).

Finally note, physical size is often used metaphorically to emphasise evaluations (e.g. “you big bully”) [34], although challengingly the bigness may be literal.

4 User Testing




Fig. 1. E-drama user interface

We conducted a two-day pilot user test with 39 secondary school students in May 2005, to try out and refine a testing methodology and with the primary aim of measuring the extent to which having EMMA as opposed to a human play a character affects levels of enjoyment, sense of engagement, etc. Fig. 1 shows a screen shot. We concealed the fact that EMMA was involved in some sessions in order to have a fair test of the difference that is made. We obtained surprisingly good results. Having a minor bit-part character called “Dave” played by EMMA made no statistically significant difference to measures of user engagement and enjoyment, or indeed to user perceptions of the worth of the contributions made by the character “Dave”. Users did comment in debriefing sessions on some utterances of Dave’s, so it was not that there was a lack of effect simply because users did not notice Dave at all. Also, the frequencies of human “Dave” and EMMA “Dave” being responded to during the improvisation (sentences of Dave’s causing a response divided by all sentences said by “Dave”) are both roughly around 30%, again suggesting that users notice Dave. Additionally, the frequency of response to other side-characters was roughly the same – “Matthew”: around 30% and “Elise”: around 35%.

Furthermore, it surprised us that no user appeared to realize that sometimes Dave was computer-controlled. We stress, however, that it is not an aim of our work to ensure that human actors do not realize this. More extensive, user testing at several

Birmingham secondary schools is being conducted at the time of writing this paper, now that we have tried out and somewhat modified the methodology.

5 Conclusion and Ongoing Work

We have implemented a limited degree of affect-detection in an automated bit-part character in an e-drama application, and fielded the character successfully in pilot user-testing. Although there is a considerable distance to go in terms of the practical affect-detection we plan to implement, the already implemented detection causes reasonably appropriate contributions by the automated character. We later intend to use affect-detection in a module for automatically generating director messages to actors. Aside from affect-sensitive directorial messages, we also intend to implement facilities for automatically generating director messages (or at least hints to the human director) when a particular character is not participating for long periods or when a character appears to be hogging the stage. In general, our work contributes to the issue of what types of automation should be included in the interactive narrative environments and how detecting affect in language can contribute to the development of believable synthetic AI characters, which contribute to the drama improvisation. Moreover, the development of affect detection provides a good test-bed for the accompanying deeper research into how affect is conveyed linguistically.

Acknowledgments. This work is supported by grants RES-328-25-0009 from the ESRC under the ESRC/EPSRC/DTI “PACCIT” programme and grant EP/C538943/1 from the EPSRC. We are grateful to our industrial partners on this grant-Hi8us Midlands Ltd, Maverick Television Ltd and BT-and to our colleagues W.H. Edmondson, S.R. Glasbey, M.G. Lee, Rachel Pilkington and Z. Wen.

References

1. Machado, I., Prada, R. and Paiva, A. 2000. Bringing Drama into a Virtual Stage. In *Proceedings of ACM Conference on Collaborative Virtual Environments*, ACM Press.
2. Picard, R.W. 2000. *Affective Computing*. The MIT Press. Cambridge MA.
3. Ortony, A., Clore, G.L. & Collins, A. 1988. *The Cognitive Structure of Emotions*. C.U.P.
4. Prendinger, H. & Ishizuka, M. 2001. Simulating Affective Communication with Animated Agents. In *Proceedings of Eighth IFIP TC.13 Conference on Human-Computer Interaction*, Tokyo, Japan, pp.182-189.
5. Wiltschko, W. R. 2003. Emotion Dialogue Simulator. eDrama learning, Inc. eDrama Front Desk.
6. Mehdi, E. J., Nico P., Julie D. & Bernard P. 2004. Modeling Character Emotion in an Interactive Virtual Environment. In *Proceedings of AISB 2004 Symposium: Motion, Emotion and Cognition*. Leeds, UK.
7. McCrae, R.R. and John, O.P. 1992. An Introduction to the Five Factor Model and Its Application. *Journal of Personality*, 60, pp.175-215.
8. Gratch, J. & Marsella, S. 2004. A Domain-Independent Framework for Modeling Emotion. *Journal of Cognitive Systems Research*. Vol 5, Issue 4, pp.269-306.

9. Egges, A., Kshirsagar, S. & Magnenat-Thalmann, N. 2003. A Model for Personality and Emotion Simulation, In *Proceedings of Knowledge-Based Intelligent Information & Engineering Systems (KES2003)*, Lecture Notes in AI. Springer-Verlag: Berlin.
10. Elliott, C., Rickel, J. & Lester, J. 1997. Integrating Affective Computing into Animated Tutoring Agents. In *Proceedings of IJCAI'97 Workshop on Intelligent Interface Agents*.
11. Mateas, M. 2002. Ph.D. Thesis. Interactive Drama, Art and Artificial Intelligence. School of Computer Science, Carnegie Mellon University.
12. Zhe, X. & Boucouvalas, A. C. 2002. Text-to-Emotion Engine for Real Time Internet Communication. In *Proceedings of International Symposium on Communication Systems, Networks and DSPs*, Staffordshire University, UK, pp.164-168.
13. Boucouvalas, A. C. 2002. Real Time Text-to-Emotion Engine for Expressive Internet Communications. In *Being There: Concepts, Effects and Measurement of User Presence in Synthetic Environments*. G. Riva, F. Davide and W. IJsselstein (eds.) pp.305-318.
14. Craggs, R. & Wood, M. 2004. A Two Dimensional Annotation Scheme for Emotion in Dialogue. In *Proceedings of AAAI Spring Symposium: Exploring Attitude and Affect in Text*.
15. Fussell, S. & Moss, M. 1998. Figurative Language in Emotional Communication. In S. R. Fussell and R. J. Kreuz (Eds.), *Social and Cognitive Approaches to Interpersonal Communication*. Lawrence Erlbaum, pp.113-142.
16. Kövecses, Z. 1998. Are There Any Emotion-Specific Metaphors? In *Speaking of Emotions: Conceptualization and Expression*. Athanasiadou, A. and Tabakowska, E. (eds.), Berlin and New York: Mouton de Gruyter, pp.127-151.
17. Barnden, J.A., Glasbey, S.R., Lee, M.G. & Wallington, A.M. 2004. Varieties and Directions of Inter-domain Influence in Metaphor. *Metaphor and Symbol*, 19(1), pp.1-30.
18. Watson, D. and Tellegen, A. 1985. Toward a Consensual Structure of Mood. *Psychological Bulletin*, 98, pp.219-235.
19. Briscoe, E. and J. Carroll. 2002. Robust Accurate Statistical Annotation of General Text. In *Proceedings of the 3rd International Conference on Language Resources and Evaluation*, Las Palmas, Gran Canaria. pp.1499-1504.
20. Metaphone Algorithm. <http://aspell.net/metaphone/>.
21. Levenshtein Distance Algorithm <http://www.merriampark.com/ld.htm>.
22. Jess, 2004. The Rule Engine for Java Platform. [Http://herzberg.ca.sandia.gov/jess/](http://herzberg.ca.sandia.gov/jess/).
23. Heise, D. R. 1965. Semantic Differential Profiles for 1,000 Most Frequent English Words. *Psychological Monographs*. 79, pp.1-31.
24. WordNet, A Lexical Database for the English Language. Version 2.1 Cognitive Science Laboratory. Princeton University.
25. Wallington, A.M., Barnden, J.A., Glasbey, S.R. & Lee, M.G. 2006. Metaphorical reasoning with an economical set of mappings. *Delta*, 22:1, pp 147-171.
26. Barnden, J.A. Forthcoming. Metaphor, Semantic Preferences and Context-sensitivity. Invited chapter for a Festschrift volume. Kluwer.
27. Moon, R 1988. *Fixed idioms and expressions in English*. Clarendon Press: Oxford, U.K.
28. Deignan, A. 2005. *Metaphor and corpus Linguistics*. John Benjamins.
29. Goatly, A. 1997. *The language of metaphors*. Routledge London and New York
30. Wallington A.M., Barnden, J.A., Barnden, M.A., Ferguson, F.J. & Glasbey, S.R. 2003. Metaphoricity Signals: A Corpus-Based Investigation. Technical Report CSRP-03-5, School of Computer Science, The University of Birmingham, U.K.
31. Wilks, Y., 1978. Making preferences more active. *Artificial Intelligence*, 10, pp. 75- 97
32. Fass, D. 1997. *Processing metaphor and metonymy*. Greenwich, Connecticut: Ablex.
33. Mason, Z.J. 2004. CorMet: a computational, corpus-based conventional metaphor extraction system. *Computational Linguistics* 30:1. pp. 23-44.
34. Sharoff, S. 2005. How to Handle Lexical Semantics in SFL: a Corpus Study of Purposes for Using Size Adjectives. *Systemic Linguistics and Corpus*. London: Continuum.