Developments in Affect Detection from Text in Open-Ended 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, United Kingdom

Abstract. We report progress[1] on adding affect-detection to an existing program for virtual dramatic improvisation, monitored by a human director. To partially automate the directors' functions, we have partially implemented the detection of emotions, etc. within users' text input. The affect-detection module has been used for the development of an automated virtual actor. The work also involves basic research into how affect is conveyed through metaphor.

1 Introduction

Improvised drama and role-play are widely used in education and training, counselling and conflict resolution. Researchers have also explored 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]). The springboard for our research was a previously existing e-drama system, *edrama*, created by Hi8us Midlands Ltd and used in schools for creative writing and teaching in various subjects. The experience suggests that e-drama (and in particular the specific system *edrama*) helps students lose their usual inhibitions, because they are not physically present on a stage and are anonymous. One main aspect of our project is the addition of types of intelligent automation to the *edrama* system.

In the *edrama* system, the virtual characters on the virtual stage are controlled by human users ("actors"), and a character's "speeches" are textual and typed in by the actor operating the character. Up to five actors and one director are involved in an e-drama session. There is a graphical interface on each actor's terminal and on the director's terminal, showing the virtual stage and the virtual characters. Characters' speeches are shown as text bubbles. Actors can choose the clothes and bodily appearance for their own characters.

The characters' visual forms have so far been static cartoon figures, and real-life photographic images have been used as the backdrops against which the characters are placed. However, we are now bringing in animated gesturing

[1] This work is supported by grant RES-328-25-0009 from the ESRC under the ESRC/EPSRC/DTI "PACCIT" programme. We are grateful to our industrial partners on this grant-Hi8us Midlands Ltd, Maverick Television Ltd and British Telecommunications plc-and to our colleagues W.H. Edmondson, S.R. Glasbey, M.G. Lee, Rachel Pilkington and Z. Wen. The work is also partially supported by EPSRC grant EP/C538943/1.

avatars and 3D computer-generated settings using technology from one of our industrial partners, BT.

Actors and the human director work through software clients connecting with the server. Clients communicate with each other by XML stream messages via the server. The server is generally at a site remote from the terminals, which may themselves be remote from each other. Terminal/server communication is over the Internet using standard browsers for the users' interaction.

The actors are given a loose scenario around which to improvise, but are at liberty to be creative. For example, one of the main scenarios we have used is a school-bullying one. In an initial phase of the e-drama session, actors are told that the schoolgirl Lisa has been bullied by her classmate Mayid. Lisa is a shy child and she is afraid of Mayid. Other characters are Lisa's mother, a friend, and a schoolteacher. Actors are expected to improvise some interesting interchanges within these parameters. It is expected that normally the Mayid character will express hostility to Lisa and that she will express fear, but actors can be creative, so that for example the Mayid actor might play him as repenting of his bullying.

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, for example, one character is not getting involved, or is dominating the improvisation.

Intervention can take a number of forms. The director may choose to send a message privately to an actor, suggesting for example that s/he interact more (or less) with another character. The director may also suggest to an actor that s/he be, for example, more aggressive or appear more upset. The director can also broadcast suggestions to all the 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 an inappropriate exchange developing between two characters.

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 a simple automated actor, EmEliza, 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 EmEliza 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 judgments* (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 can be useful. We do not purport to be able to make EmEliza detect all types of affect under all ways affect can be expressed or implied, or to do it with a high degree of reliability. 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. Indeed, 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 in part to reason about emotions and to produce believable emotional expression. Wiltschko's *eDrama Front Desk* [5] is designed as 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 their approach for the generation of 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. Although *Facade* [11] included shallow natural language processing for characters' open-ended utterances, 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. There has been relevant 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 how scientifically accurate those views are. 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 has been used to deal with abbreviations (e.g. 'im (I am)' and 'c u (see you)'). Including most abbreviations used in Internet chat rooms and textese, it can handle most abbreviation in users' input. Especially we also deal with abbreviations such as numbers embedded within words (e.g., "l8r" for later) using the lookup table. We handle the ambiguity of, for example,"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. Such simple processing inevitably leads to errors in some cases, but in evaluations using examples in previous transcripts we have obtained 85.7% accuracy, which is adequate currently.

Letter repetition comes in two flavours. One is repetition added to ordinary words (e.g. 'yessss', 'seeeee') and the other is repetition added to interjections or onomatopoeic elements (e.g. 'grrrrr', 'agggghhh'). The iconic use of word length here (i.e., written word length corresponding roughly to imagined sound length) normally implies strong affective states in the characters' input. Usefully, adding letters does not change the pronunciation a great deal. 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 well 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 preprocessing, we turn to the core detection of affect in users' input.

2.2 Affect Detection by Pattern Matching

In an initial stage of our work, affect detection was based 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 the 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 in speeches. 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 based on a school bullying scenario. 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 EmEliza's different roles in specific scenarios.

A rule-based Java framework called Jess [22] is used to implement the pattern-matching rules in EmEliza. When the character Mayid says "Lisa, you Pizza Face! You smell", EmEliza 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 EmEliza 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 EmEliza tries to respond, this example rule will be fired. EmEliza 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 emotion and exclamation marks or capitalisation 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

One useful pointer to affect is the use of imperative mood, especially when used without softeners such as 'please' or 'would you'. Strong emotions and/or rude attitudes are often expressed in this case. There are common imperative phrases we deal with explicitly, such as "shut up" and "mind your own business". They usually indicate strong negative emotions. But the phenomenon is more general.

Detecting imperatives accurately in general 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 at least some of the difficulties.

The Rasp parser recognises some types of imperatives directly. Unfortunately, the grammar of the 2002 version of the Rasp parser that we have used does not deal properly with certain imperatives (John Carroll, p.c), which 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 statement sentences or in the latter example an interrogative. Therefore, further analysis is needed to detect imperatives, by additional processing applied to the possibly-incorrect syntactic trees produced by Rasp. This includes consideration of the nature of the sentence subject, what form of the verb is used and whether negation is present. We mention one case of special interest, as follows. It brings in a need for semantic and pragmatic processing.

When a sentence involves an explicit subject and a verb for which there is no difference at all between the base form and the past tense form, then ambiguity arises between imperative and declarative (e.g. "Lisa hit me"). There is an important special case of this ambiguity. If the object of the verb is 'me', then 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 to the others, if the evaluation value is negative for such a verb, then this sentence is probably not 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.

Aside from imperatives, we have also worked on implementing simple types of semantic extraction of affect using affect dictionaries and electronic thesauri, such as WordNet [24]. The way we are currently using WordNet is briefly as follows.

2.4 Using WordNet for a First Person Case

As we mentioned earlier, use of 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 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 EmEliza'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

EmEliza normally responds to, on average, every Nth speech by another character in the e-drama session, where N is a changeable parameter (currently usually set to 3). However, it also responds when EmEliza's character's name is mentioned, and makes no response if it cannot detect anything useful in the utterance it is responding to. The one-in-N average is achieved by sampling a random variable every time another character says something. As a future development we plan to have N dynamically adjustable according to how confident EmEliza is about what it has discerned in the utterance at hand.

EmEliza makes a random response from several stored response candidates that are suitable for the affective quality it has discerned in the utterance it is responding to. For example,

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

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

In addition, EmEliza sometimes takes 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 multi-word phrases or the original words, which translate the abbreviations, and the correct spelling ones in the users' input, this helps to obscure the fact that part of EmEliza's response is only a reflection. For example,

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

<EmEliza character> Why are you my best mate?

On the other hand, we also purposely put abbreviations and other phenomena used in the Internet chat rooms in EmEliza's responses in order to make EmEliza 'behave' like a school kid. For example,

<Another character> You get lost!

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

<Another character> He didn't do it.

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

The reader should bear in mind that EmEliza'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.

3 Affect via Metaphor

The direct metaphorical description of emotional states is common and has been extensively studied [15]. Examples are "He nearly exploded" and "Joy ran through me." We say that such descriptions are "direct" because they are directly about emotional states, even though in many cases no emotional state is named. But affect is often conveyed indirectly via metaphor, as in "His room is a cess-pit", where affect associated with a source item (cess-pit) gets carried over to the corresponding target item (the room). Here the sentence is not about an affective quality but nevertheless an affective quality is conveyed.

In our research on metaphor (see, e.g., [17, 25]) we are concerned with metaphor in general and are in particular interested in both of the types of affective metaphor in the previous paragraph. We are bringing this metaphor research to bear upon the e-drama application, and using the application as a useful source of theoretical inspiration.

Our intended approach to metaphor handling in the EmEliza affect-detection module is partly to look for stock metaphorical phraseology and straightforward variants of it, and partly to use a simple version of the more open-ended, reasoning-based techniques taken from the ATT-Meta project on metaphor processing [17]. As an example of stock phrase handling, insults in e-drama are often metaphorical, especially the case of animal insults ("Don't you dare call her a pig, you dog"). Simple pattern-matching rules are currently used in EmEliza to deal with such insulting language.

In most discourse, metaphorical phraseology tends to be of conventional form, the extreme being stock phrases such as"sit on the fence". Such phrases can be stored in a lexicon and directly recognized (although there is still the problem that they might be used literally in some contexts). However, it is quite common also to get variations of stock phraseology, of a nature that defeats any simple lexicon-based approach and raises the need for a certain amount of knowledge-based reasoning, as in the possible variant "sit on a very shaky fence". Such cases would benefit from the reasoning capabilities of ATT-Meta. The phenomenon is discussed in more detail in [25].

A particular feature of the ATT-Meta approach is the identification of certain very general ways in which qualities are transmitted in metaphor from information about the source item (e.g., cess-pit) to the target-item (e.g., a room). No matter what the particular correspondences are involved in a metaphor, it appears that certain types of quality map across from source to target (defeasibly-other knowledge may defeat the effect). For instance, qualitative temporal relationships map across. But another important case is precisely the mapping across of affective qualities. The cess-pit example provides a simple case of this, assuming that cess-pits have by default a negative affective connotation.

Another, related, phenomenon is that physical size is often metaphorically used to emphasize evaluations, as in "you are a big bully", "you're a big idiot", and "you're just a little bully." The bigness is sometimes literal as well. "Big bully" expresses strong disapproval [26] and "little bully" can express contempt, although note that "little" often also conveys positive affect such as sympathy or affection. Such cases pose interesting challenges to our ATT-Meta based approach.

We have encountered creative uses of metaphor in e-drama. For example, in the school-bullying scenario, Mayid has already insulted Lisa by calling her a 'pizza' (short for 'pizza-face'). Then in one improvisation Mayid said "I'll knock your topping off, Lisa" - a theoretically intriguing spontaneous creative elaboration of the 'pizza' metaphor.

4 User Testing

We conducted a two-day pilot user test with 39 secondary school students in May 2005, in order to try out and refine a testing methodology. We concealed the fact that EmEliza was involved in some sessions. We got surprisingly good results regarding EmEliza (even though it was not then using the more sophisticated techniques in the present paper) judging by feedback from the school students in the group debriefing sessions after the e-drama sessions. In particular, nobody found out that (sometimes) one bit-part character was computer-controlled, despite the fact that the character's interventions did have an effect on various measures assessed via questionnaire (statistical analysis still ongoing). Further user testing with groups of secondary school students at several Birmingham schools will take place in early 2006. The main manipulation will again be the inclusion or otherwise of an EmEliza-controlled bit-part character. We will be statistically analyzing the effect of this manipulation on the users' overall enjoyment, sense of engagement, nervousness, etc.

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 that we plan to implement, the already implemented detection is able to cause reasonably appropriate contributions by the automated character.

We also intend to use the affect-detection in a module for automatically generating director messages to actors. We envisage the human director specifying certain types and levels of emotion, etc. that particular characters are expected to display, or having these specified as part of the scenario, so that the affect-detection module can look out for them. A difficulty is allowing for creative improvisation that introduces an unexpected but nevertheless valuable affective profile. We can expect to be able to do only a limited amount about this in automation, but at the very least the human director could dynamically adjust the expected affect parameters if an unexpected but desirable profile arises so that the change can affect the automated directorial functions.

Aside from affect-sensitive directorial messages, we are in the process of implementing facilities for automatically generating director messages (or at least hints to the human director about what messages to send) when a particular character is not participating for long periods or when a character appears to be hogging the stage.

References

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

8. Gratch, J. and Marsella, S.: A Domain-Independent Framework for Modeling Emotion. *Journal of Cognitive Systems Research.* Vol 5, Issue 4, (2004) 269-306
9. Egges, A., Kshirsagar, S. and Magnenat-Thalmann, N.: A Model for Personality and Emotion Simulation. In *Proceedings of Knowledge-Based Intelligent Information and Engineering Systems*, Lecture Notes in AI. Springer-Verlag: Berlin. (2003)
10. Elliott, C., Rickel, J. and Lester, J.: Integrating Affective Computing into Animated Tutoring Agents. In *Proceedings of IJCAI'97 Workshop on Intelligent Interface Agents.* (1997)
11. Mateas, M.: Ph.D. Thesis. Interactive Drama, Art and Artificial Intelligence. School of Computer Science, Carnegie Mellon University.(2002)
12. Zhe, X. and Boucouvalas, A. C.: Text-to-Emotion Engine for Real Time Internet Communication. In *Proceedings of International Symposium on Communication Systems, Networks and DSPs.* Staffordshire University, UK, (2002) 164-168
13. Boucouvalas, A. C.: 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. IJsselsteijn (eds.) (2002) 305-318
14. Craggs, R. and Wood. M.: A Two Dimensional Annotation Scheme for Emotion in Dialogue. In *Proceedings of AAAI Spring Symposium: Exploring Attitude and Affect in Text.* (2004)
15. Fussell, S. and Moss, M.: Figurative Language in Descriptions of Emotional States. In S. R. Fussell and R. J. Kreuz (Eds.), *Social and Cognitive Approaches to Interpersonal Communication.* Lawrence Erlbaum. (1998)
16. Kovecses, Z.: 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. (1998) 127-151
17. Barnden, J.A., Glasbey, S.R., Lee, M.G. and Wallington, A.M.: Varieties and Directions of Inter-domain Influence in Metaphor. *Metaphor and Symbol*, **19(1)** (2004) 1-30
18. Watson, D. and Tellegen, A.: Toward a Consensual Structure of Mood. *Psychological Bulletin*, **98** (1985) 219-235
19. Briscoe, E. and Carroll. J.: Robust Accurate Statistical Annotation of General Text. In *Proceedings of the 3rd International Conference on Language Resources and Evaluation*, Las Palmas, Gran Canaria. (2002) 1499-1504
20. Metaphone Algorithm. http://aspell.net/metaphone/. (2005)
21. Levenshtein Distance Algorithm http://www.merriampark.com/ld.htm. (2005)
22. Jess. The Rule Engine for Java Platform. http://herzberg.ca.sandia.gov/jess/. (2004)
23. Heise, D. R.: Semantic Differential Profiles for 1,000 Most Frequent English Words. *Psychological Monographs.* **79** (1965) 1-31
24. WordNet, A Lexical Database for the English Language. Version 2.1 Cognitive Science Laboratory. Princeton University. (2006)
25. Barnden, J.A.: Metaphor, Semantic Preferences and Context-sensitivity. Invited chapter for a Festschrifft volume. Kluwer. (to appear)
26. Sharoff, S.: How to Handle Lexical Semantics in SFL: a Corpus Study of Purposes for Using Size Adjectives. *Systemic Linguistics and Corpus.* London: Continuum. (2005)

This article was processed using the LaTeX macro package with LLNCS style