eDrama: Facilitating Online Role-play using Emotionally Expressive Characters

Kulwant Dhaliwal*, Marco Gillies[†], John O'Connor*Amanda Oldroyd[‡], Dale Robertson[‡], Li Zhang[☼]
*Hi8us Midlands, Unit F1, The Arch, 48-52 Floodgate Street, Birmingham, B5 5SL
[†]Department of Computer Science, University College London, Malet Place, London WC1E 6BT
[‡]BT Group CTO, Adastral Park, Ipswich, Suffolk, IP5 3RE
[☼]School of Computing and Technology, University of East London, Docklands Campus, 4-6 University Way, London E16 2RD

1 Introduction

This paper describes the results of a user study of a multi-user role-playing environment '*edrama*', which enables groups of people to converse online, in scenario driven virtual environments. Hi8us' *edrama* system is a 2D graphical environment in which users are represented by static cartoon like avatars.

An application has been developed to enable the integration of the existing *edrama* tool with several new software components to support avatars with emotionally expressive behaviours, rendered in a 3D environment.

In this paper we describe a user trial that demonstrates that the changes made improve the quality of social interaction and users' sense of presence.

2 Overview of *edrama*

In 1999, Hi8us Midlands began developing *edrama* — online multi-user role-play software that could be used for education or entertainment. This first major *edrama* project received substantial support from NESTA, the National Endowment for Science Technology and the Arts. *edrama* became one of Nesta's flagship projects.. In this software young people could interact online in a 2D flash based interface with others under the guidance of a director. The interface incorporated 2D static avatars and a text chat interface, with different photographic backgrounds as scenes to set the role-play. Over the years, edrama has been further developed by Hi8us Midlands and adapted for the delivery of commissions for a range of uses, such as Careers Advice and Creative Writing. The Dream Factory was a version of *edrama* developed for the University for Industry and piloted and tested Connexions-Direct advisers and young people at Skill City in Salford. The software continues to be developed as a 2D application, to be used online

The benefits of Hi8us' *edrama* 2D software include:

- Its use of drama to deliver almost any type of training, which is engaging and entertaining
- It's collaborative and multi-user, allowing people to learn together remotely, cutting out any geographical, social and cultural barriers
- One of the advantages of the 2D version is that anyone with an internet connection, on any platform, can access it – once loaded, it can even be used via a 56k modem and has been trialled in situations such as these
- Role-play in person is often an area that many shy away from due to inhibitions around performing, but *edrama* allows users to remain anonymous, allowing them to express themselves with out being seen — this is particularly useful with young people who may be afraid of expressing their views in front of their peers
- *edrama* can easily be customised, which is a real benefit for trainers wanting to create role-plays to make them more specifically relevant for their purposes or, change the scene backdrops available – in a few moments a photo taken on a digital camera can be transferred into the tool.
- edrama is chat with purpose, it builds on a popular pastime activity amongst users young and old, but gives a framework that allows it to be purposeful activity
- Facilitation of the role-plays is a crucial aspect of the tool, this ensures that the 'chat' is purposeful and assists users to respond to the given situations
- The text from each role-play is automatically saved, which means that there is a record of every session, which can be used for assessment purposes

Although the 2D version of *edrama* has been successfully used in a number of situations and continues to provide that capacity with further opportunities to implement it in the pipeline, it has the potential to benefit from additional features. One of the main reasons that Hi8us undertook the collaborative eDrama project, was to explore these areas and see how they might benefit the user experience and delivery of the tool. For example, can edrama, delivered as a 3D environment, allow users to be more thoroughly immersed? Or, could the role-plays be automatically facilitated and be made more engaging by the use of AI agents?

This paper describes an alternative version of the *edrama* software developed in collaboration with Hi8us Midlands, Maverick TV, Birmingham University and BT with the support of the PACCIT programme (People at the Centre of Communication and Information Technologies). Our collaboration aims to enrich the user-experience with emotionally responsive characters, including additional non-human characters within a 3D application. The addition of 3D capabilities include character and background scene rendering and enables real-time processing of animation to visually update the current emotional state of every character on screen.

In the following sections we will describe the three main components integrated into *edrama*, which support the addition of expressive characters. We describe the prototype actor application supporting 3D backgrounds and avatars and integration with existing *edrama* functionality. We then cover the two components that create the expressive characters – the improvisational AI bit-part character and Demeanour. We go on to outline the user study, including the scenarios used, the experiment setup and results.

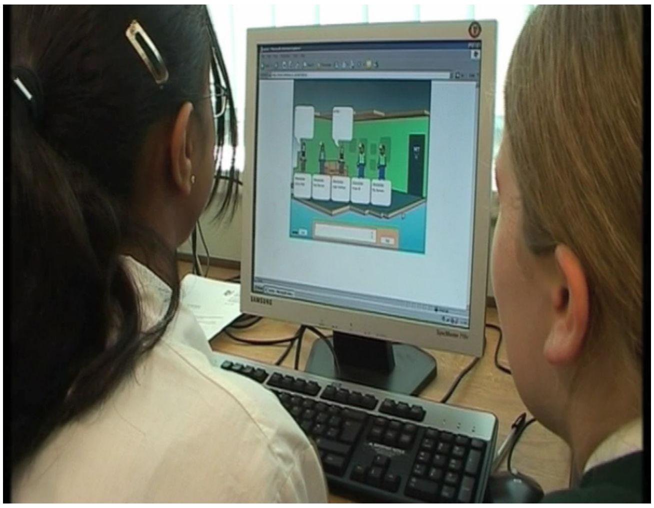

Figure 1 Children using Hi8us' 2D version of *edrama*

3 *edrama* application

The *edrama* software consists of two main user interfaces, an 'actor' client application that is used by the actors, and a 'director' client application which is available solely to the director. The director interface remains largely unchanged. It is a web-based interface which incorporates a number of tools to start role-play sessions, view the scene and avatars (in 2D), and monitor the conversation. The director can start, stop and change background scenes, and to talk to one or all of the participants using text chat. In contrast, the actor client has undergone significant developments to support the real-time rendering of expressive characters and is the focus of the user study.

The 3D version of the *edrama* actor client is an MS Windows based application written using MFC. The application consists of two child windows; one houses the Flash* Player ActiveX control to enable Flash movies to be played within that window, the remaining is a TARA enabled window that displays the 3D visuals. Hi8us' *edrama* is a web-browser hosted Flash movie. This is noted because though the structure of the *edrama* client may have changed user interaction is still controlled through a Flash movie interface for consistency with Hi8us' versions.

TARA is an SDK developed by BT used to create real-time 3D enabled applications. The SDK provides a set of extensible components that are used to render geometry and effects using MS DirectX. The TARA SDK allows its core components to be replaced for more functional components tailored to meet a specific need. This provides a mechanism in which to integrate new technologies into TARA enabled applications, without the components of the SDK needing to know about them; in this case the Demeanour framework (Gillies et al, 2006). The creation of an alternative 3D system was an attempt to enrich the environment provided by eDrama to reinforce the emotional content of the role-play..

The flash movie developed for TARA, is the user interface that controls the flow of the application. It is a simplified version of the 2D web-based interface so that a move from one to another would require no learning on the part of the user; this also means that this prototype is compatible with the current Hi8us version therefore the two can be used in parallel. The flash movie is the client to the external server passing state related messages for each *edrama* client and capturing messages broadcast by the server. The flash interface communicates with the *edrama* application through a socket maintained by the application. Messages broadcast by the server are passed through this socket to be processed. The communication between interface and application is one way only, from interface to application as the application is responsible for reflecting the current state based upon users interactions.

To take part in an *edrama* session a user runs the actor client locally on their desktop. The user must first login, to select an available session and character to play. The login screen enables each actor to login with a unique ID. This ID can be anonymous – a numerical identifier for example or can be a username. In all cases this login does not require any personal details and is not referred to in the role-play. The login interface presents a number of options, including key stage (if required) , the role-play option and the characters available in the session. A total of 5 characters are available in each role-play.

Customisation of the character takes place in a virtual 'dressing room', available after login. This includes scenario details and customisation tools. The 3D window includes an interactive text panel which displays background information about the scenario and selected character, and also renders a default avatar in the 3D dressing room, animated with general waiting poses. Actual customisation is through the e-fit tool in the flash interface, which provides click through image based selection of gender, head, torso and leg options. The results are displayed in real time in the 3D dressing room window.

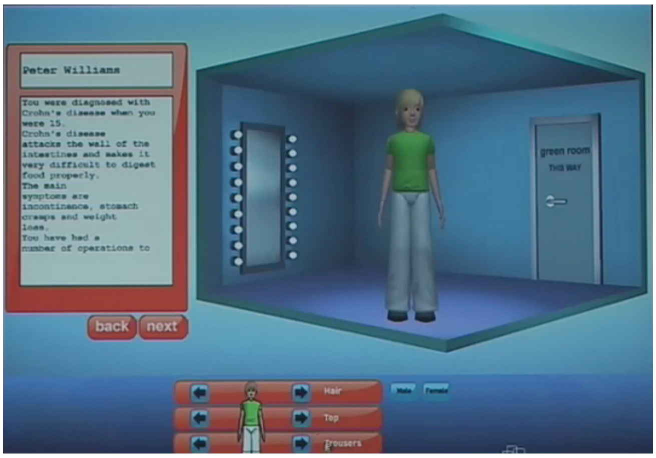

Figure 2 Avatar customisation

From this point the user moves into the multi-user sections. The first of these is the 'green room', which is a warm up space to meet other actors and the director. The 3D window displays the green room scenery and all logged in avatars. Text chat is displayed in speech bubbles above the avatars' heads. Text is input via the flash panel.

The director will enable the stage once all the characters have appeared in the green room and warmed up. The director signals the start of the role-play using an 'Action' command, which warns the actors of the scene change. A background scene from a library of options is displayed on all clients. This maybe updated at anytime during the session. The role-play is ended using a 'Cut' command from the director – at which point the application will close down all actor clients.

In both the green room and stage environments, each actor is given a set position on screen, resulting in a semi-circle of characters facing camera (the user's viewpoint). In this case the actor can see their own avatar in the 3rd person as part of the avatar group.

Figure 3 Four actors in the green room

This tableaux format provides the user a view of the whole role-play and avatar positions are consistent on each actor client. There is no ability to navigate the scene. This means that time is not taken up with actors trying to negotiate places on screen, or to have to arrange themselves so they can all be seen clearly. Additionally the actors can concentrate completely on talking to each other and watching the unfolding scenario on screen.

When the director speaks to the group or individual avatars a 2D director image overlays the window and text appears in a speech bubble. In this way the director can appear to the group or to single clients and give directions to assist the role-play.

When an actor types in text in the chat panel, the text appears in bubbles above the avatar. Each character is animated according to its emotional profile and to the text input of users during the session.

4 Emotionally expressive characters

Each scenario has a written description, or profile, of 5 characters who are able to participate. There is usually a main character or protagonist, who faces a conflict or issue, this character will have a counterpart who is the antagonist and takes an opposing view. The remaining characters will have specific relationships to these characters (parent, friend, enemy). This information is provided in the character background information. In many scenarios the basic character profiles have a similar pattern to provide the basis of a productive role-play. In Hi8us' versions of edrama this information can only inform the performance of the actor engaged in the role-play. In the 3D version described here it becomes influential in how the avatars are animated on screen.

Using a combination of character profiles and detected affective states from user's text it is possible to animate each character with expressive behaviour, without any direct user intervention via the *edrama* interface. This employs a combination of two technologies, affect detection in open-ended improvisational text (Zhang et al. 2006) and Demeanour framework (Gillies et al. 2006)

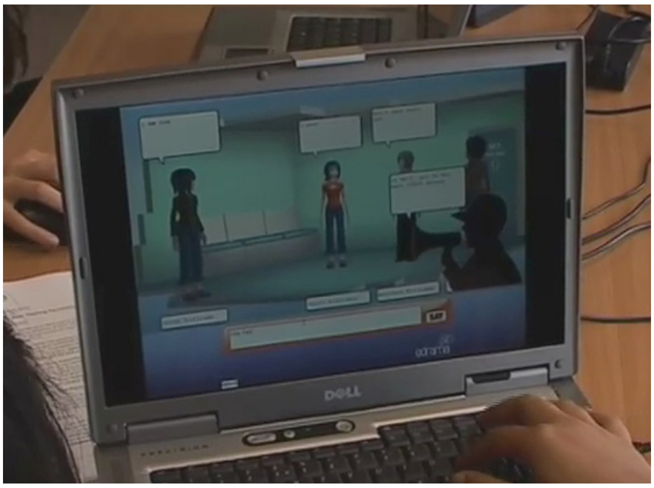

Figure 4 Director interacts with the actors

Affect detection in open-ended improvisational text

In *edrama*, the actors (users) are given a scenario within which to improvise, but are at liberty to be creative. There is also a human director, who constantly monitors the unfolding drama and can intervene by, for example, sending messages to actors, or by introducing and controlling a minor 'bit-part' character to interact with the main characters. 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. Additionally, the Director role was originally designed to be undertaken by Teachers, but it is now easily performed by pupils as well as teachers and this works quite well. It has also successfully been delivered by Careers Advisers, who have received no more than 30 minutes of training, to successfully perform the director role. However, within all sectors, commercial and otherwise, the need to cut costs in terms of staff time to deliver services is of great importance.

One research aim is thus partially to automate the directorial functions, which importantly involve affect detection. For instance, a director may intervene when emotions expressed or discussed by characters are not as expected. Hence we have developed an affect-detection module. The module identifies affect in characters' text input, and makes appropriate responses to help stimulate the improvisation. Within affect we include: basic and complex *emotions* such as anger and embarrassment; *meta-emotions* such as desiring to overcome anxiety; *moods* such as hostility; and *value judgments* (of goodness, etc.). Although merely detecting affect is limited compared to extracting full meaning, this is often enough for stimulating improvisation. The results of this affective analysis are then used to: (a) control an automated improvisational AI actor – EMMA (emotion, metaphor and affect) that operates a bit-part character in the improvisation; (b) drive the animations of the avatars in the user interface so that they react bodily in ways that is consistent with the affect that they are expressing, for instance by changing posture or facial expressions. The response generation component of EMMA uses this interpretation to build its behaviour driven mainly by EMMA's role in the improvisation and the affect expressed in the statement to which it is responding. The intention of EMMA's response is to hopefully stimulate the improvisation.

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* (Mateas, 2002) 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 (2002) demonstrated an emotion extraction module embedded in an Internet chatting environment. 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 (Craggs & Wood, 2004).

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 (Fussell & Moss, 1998), and to push forward the theoretical study of such language, as part of our research on metaphor generally (see, e.g. Barnden et al., 2004).

Our affect detection module

The language in the textual 'speeches' created in *edrama* 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 abbreviations, misspellings, slang, use of upper case and special punctuation (such as repeated exclamation marks) for affective emphasis, repetition of letters, syllables or words for emphasis, and open-ended interjective and onomatopoeic elements such as "hm", "ow" and "grrrr". To deal with the misspellings, abbreviations, letter repetitions, interjections and onomatopoeia, several types of pre-processing occur before the main aspects of detection of affect. We have reported our work on pre-processing modules to deal with these language phenomena in detail in Zhang et al. (2006).

Now we briefly introduce our work on the core aspects of affect detection. 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. Expression of the imperative mood in English is surprisingly various and ambiguity-prone, as illustrated below. We have used the syntactic output from the *Rasp* parser (Briscoe & Carroll, 2002) and semantic information in the form of the semantic profiles for the 1,000 most frequently used English words (Heise, 1965) to deal with certain types of imperatives. Briefly, the grammar of the 2002 version of the *Rasp* parser that we have used incorrectly recognised certain imperatives (such as "you shut up", "Dave bring me the menu" etc) as declaratives. We have made further analysis of the syntactic trees produced by *Rasp* by considering of the nature of the sentence subject, the form of the verb used, etc, in order to detect imperatives. We have also made an effort to deal with one special case of ambiguities: a subject + a verb (for which there is no difference at all between the base form and the past tense form) + "me" (e.g. 'Lisa hit/hurt me'.). The semantic information of the verb obtained by using Heise's (1965) semantic profiles, the conversation logs and other indicators implying imperatives help to find out if the input is an imperative or not.

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 specific words or sets of specific alternative words. This continues to be a core aspect of our system but we have now added robust parsing and some semantic analysis, including but going beyond the handling of imperatives discussed above.

A rule-based Java framework called Jess is used to implement the pattern/template-matching rules in EMMA allowing the system to cope with more general wording. 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 characters' utterances. The rules conjecture the character's emotions, evaluation dimension (negative or positive), politeness (rude or polite) and what response EMMA should make. 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 the nature of specific scenarios.

However, it lacks other types of generality and can be fooled when the phrases are suitably embedded as subcomponents of other grammatical structures. In order to go beyond certain such limitations, sentence type information obtained from the *Rasp* parser has also been adopted in the pattern-matching rules. This information not only helps EMMA to detect affective states in the user's input (see the above discussion of imperatives), and to decide if the detected affective states should be counted (e.g. affects detected from conditional sentences won't be valued), but also helps EMMA to make appropriate responses. Additionally, the sentence type information can also help to avoid the activation of multiple rules, which could lead to multiple detected affect results for one user's input. Mostly, it will help to activate only the most suitable rule to obtain the speaker's affective state and EMMA's response to the human character.

Additionally, a reasonably good indicator that an inner state is being described is the use of 'I' (see also Craggs & Wood (2004)), especially in combination with the present or future tense (e.g. 'I'll scream', 'I hate/like you', and 'I need your help'). We especially process 'the first-person with a present-tense verb' statements using WordNet. When we fail to obtain the speaker's affective state in the current input by using Rasp and pattern matching, WordNet is used to find the synonyms of the original verb in the user's input. These synonyms are then refined by using Heise's (1965) semantic profiles in order to obtain a subset of close synonyms. The newly composed sentences with the verbs in the subset respectively replacing the original verb, have extended the matching possibilities in the pattern-matching rules to obtain user's affective state in the current input.

After the automatic detection of users' affective states, EMMA needs to make responses in her role to the human characters during the improvisation. We have also created responding regimes for the EMMA character. Most importantly, EMMA can adjust its response likelihood according to how confident EMMA is about what it has discerned in the utterance at hand.

Details of the work reported in this section can be found in Zhang et al. (2006). The brief summaries here of our previous implementations and their capabilities aim to remind readers.

The detected affective states in the user's text input and EMMA's responses to other characters have been encoded in an xml stream, which is sent to the server by EMMA. Then the server broadcasts the xml stream to all the clients so that the detected affective states information can be picked up by the animation engine to contribute to the production of 3D gestures and postures for the avatars. Now we will discuss the generation of emotional believable animation in detail in the following section.

The users avatars and emotional animation

The topics discussed in the *edrama* scenarios are often highly emotionally charged and this is reflected in the animation of the characters. Each participant in *edrama* has their own animated graphical character (avatar). In order for the characters to enhance the interaction the characters all have emotionally expressive animations. [sic] Garau *et al.* (2001) point out that avatars that do not exhibit appropriate emotional expression during emotionally charged conversation can be detrimental to an interaction. The problem with animated avatars is that they can be very complex to use if users have to directly control the avatars animation. Vilhjálmsson and Cassell (1998, 1999) have shown that users find controlling animated avatars difficult and their experience and interaction is improved if they use an avatar whose behaviour is controlled autonomously. We therefore have an autonomous model of affective animation for our avatars based on the affective states detected in users' text input. These detected affective states control the animation of the user avatars using Demeanour expressive animation framework (Gillies et al., 2006).

Demeanour makes it possible for our characters to express the affective states detected by EMMA. When EMMA detect an affective state in a user's text input, this is passed to the demeanour system attached to this user's character and a suitable emotional animation is produced. The animation system is based around a set of short animation clips, each of which is labelled with one or more affective states. Each clip only affects an individual part of the body (torso, legs, arms) and thus several clips can be easily combined at the same time. When a new affective state is received a new set of clips is chosen at random from the clips labelled with that state and these new clips are combined together to produce a new animation. Every few seconds the set of clips used is varied at random to produce a new animation, but one which has the same affective state as before. This allows us to produce varied behaviour over long time periods. The animation system also implements affective decay. Any affective state will eventually revert to a neutral state if it is not replaced by a new one.

Another feature of the animation system is that characters can produce affective responses to the states of other characters. If a character produces a strong affective state then other characters will also produce a milder response. Each character has a profile which specifies how it responds to the behaviour of each other character. This makes it possible to implement different responses for different characters. For example, two characters with a positive relationship may empathise with each other, when one is unhappy so is the other. On the other hand if two characters have a negative relationship then one might gloat at the other's unhappiness, and therefore display happiness.

Demeanour generates a number of output affective states, which are used to select the animation clips. Each output state is a weighted sum of a number of input factors. The primary input factor is the affective state as detected from the input text, this always has weight 1. The inputs also include the states of other characters, with lower weights. So for example, the output state "happiness" depends on the input "happiness", but also on the "happiness" and "sadness" values of other characters. The weights of the other characters states are contained in the character's profile. The profile consists of a separate set of weights for each other character in the scenario. This makes it possible to respond differently to each character. For example, if two characters, A and B have a poor relationship and A is angry with B, B might respond by being angry back. On the other hand if B's parents were angry then B might be sad or submissive. At any time each character has a single focus of social attention (which is itself another character), determined by the character's direction of gaze (which is itself determined by an animated gaze model). In order to generate animation the first step is to update the input states based on any text typed in. Next the states of the current focus of attention are fetched. These are multiplied by the weights of given by the profile specific to the focus of attention and added to the input emotion to produce the output state.

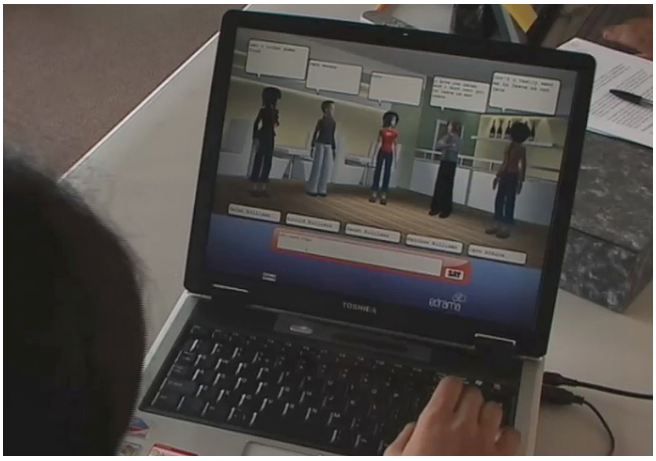

Figure 5 Children using the 3D version of *edrama*

Table 2. The Categories of questions used for the questionnaires

Category	Number of questions	Example Questions
Enjoyment	12	“ How much did you enjoy the roleplays?”
Difficulty	14	“I needed help to use edrama”
Presence	2	“I forgot I was at school when I was doing the role-play”
Co-Presence	8	“ Did you feel close to the group online?”
Quality of Social Interaction	6	“ Did you get to have your say?”
Own Avatar Appearance	3	“ I wanted the Avatar to look more like me”
Own Avatar Behaviour	5	“ My Avatar was expressive”
Other Avatars	10	“I was paying attention to other people’s Avatars”

5 The User Study

Our user study involved two trials of the prototype 3D *edrama* application. These were completed in July and October 2006, as new animation capabilities were added to the prototype system. This section describes the scenarios, the user study and results.

The scenarios

Three scenarios were used in the user testing; the first was entitled ‘Big Night Out’ which was delivered in the 2D version and served as a warm up; the other two were homophobic bullying and Crohn’s disease both of which are described below. In each case, introductory video produced by Maverick TV were shown to the trialists. These videos were case studies of both subject areas and featured interviews with either victims of bullying or sufferers of Crohn’s disease. This is additional information to help participants identify with the sensitive issues being explored in the scenarios. In these scenarios, Mr Dhanda (Homophobic Bullying) and Dave (Crohn’s Disease) are AI characters driven by EMMA.

Homophobic Bullying

In this scenario the character Dean (16 years old), captain of the footbal team, is confused about his sexuality. He has ended a relationship with a girlfriend because he thinks he may be gay and has told her this in confidence. Tiffany (ex-girlfriend) has told the whole school and now Dean is being bullied and concerned that his team mates on the football team will react badly. He thinks he may have to leave the team. The other characters are; Tiffany who is the ring leader of the bullying, and wants Dean to leave the footbal team, Rob (dean’s younger brother) and wants Dean to say he is not gay to stop the bullying, Lea (Dean’s older sister) who wants Dean to be proud of who he is and ignore the bullying, and Mr Dhanda (PE Teacher) who needs to confront Tiffany and stop the bullying.

Crohn’s Disease:

In this scenario the character Peter has had Crohn’s deisease since the age of 15. Crohn’s disease attacks the wall of the intestines and makes it very difficult to digest food properly. The character has the option to undergo surgery (ileostomy) which will have a major impact on his life. The task of the role-play is to discuss the pros and cons with friends and family and decided whether he should have the operation. The other characters are; Mum, who wants Peter to have the operation, Matthew (older brother) who is against the operation, Dad who is not able to face the situation, and David (the best friend) who mediates the discussion. The setting is a night out for an evening meal.

Procedures

There were 3 conditions in the user study:

1. Hi8us’ 2D version of *edrama* with no animation or affect detection
2. The 3D version of *edrama* with the bit part character but limited animation
3. The 3D version with the bit part character and full animation

In the version with limited animation the animations only occurred when an emotion was detected by the emotion detection system and there was only one animation per emotion. In the full animation condition animations were constantly being played and there were a variety of possible animations for each emotion.
The comparison between the 2D version and the 3D versions was performed within subjects while the comparison between the two 3D conditions was performed between subjects. The participants were therefore divided into two groups as show in the table.

Table 1. Experimental Conditions

Group	Condition 1	Condition 2
A	2D	3D with limited animation
B	2D	3D with full animation

The two groups were tested in different sessions. Group A used the 2D and 3D versions on different days while group B used them on the same day.
There were 10 participants per group. The participants were all female aged between 13 and 14 and pupils at Swanshurst School, a Specialist Science College in Billesley, Birmingham.
The participants were randomly assigned into groups of 4 and given a scenario and character. None of the participants knew who

the other members of their group were. However, due to the proximity of the terminals, sometimes they were able to establish identities of fellow participants.
The participants were then asked to role play using the *edrama* system for 10-15 minutes per session, they undertook 3 sessions in the first trial but only 2 sessions in the second one. They had less time to undertake the 2D session in the second trial, due to the technical difficulties.

Results

Participants were asked to complete a questionnaire about their experience with the 2D *edrama* before using the 3D version and a second one after using the 3D version. The two questionnaires were mostly identical, but some minor changes were made to the questions to make them applicable, and 7 questions were added that were not applicable to the 2D version. The First questionnaire had 71 questions and the second had 78. All of the questions were 7 point Likert like scales.
The questions were divided into 7 categories, shown in table 2. For each participant the mean was taken of their answers to for the questions in each category and this was used as their score for that category. The mean was then taken for each condition for each category.
The first comparison was between the 2D condition and the two 3D conditions:

Table 2 Comparison of 2D and 3D conditions

Category	2D Mean	3D Mean	t-value
Group A			
enjoyment	5.075	4.875	-0.418
difficulty	2.042	2.07	0.087
presence	4.9	3.6	-3.00
co-presence	3.5	3.5	0
Social Dynamics	4.05	4.95	2.347
Avatar Appearance	3.266	3.7	0.979
Avatar Behaviour	3.3	2.9	-0.669
Other Avatars	4.483	4.03	-1.33
Group B			
enjoyment	4.5	5.516	1.53
difficulty	3.278	2.528	-1.494
presence	4.8	5.35	0.992
co-presence	3.614	4.137	1.016
Social Dynamics	3.633	5	2.306
Avatar Appearance	3.76	3.033	-1.568
Avatar Behaviour	3.85	4.96	1.863
Other Avatars	4.5	5.35	2.072

The main significant result that was obtained consistently between the groups was that the quality of social interaction improved with the 3D condition. Interestingly the participants reported evaluation of the avatars was not significantly different between conditions, except for the case of the other avatars in the full animation condition. This might be because the participants did the first questionnaire before doing the 3D version and so were not directly comparing the avatars in the two systems.
The second comparison was between the 3D conditions with limited and full animation:

Table 3 Comparison of limited animation and full animation conditions

Category	Group A	Group B	t-value
enjoyment	4.875	5.516	1.324
difficulty	2.071	2.528	1.186
presence	3.6	5.35	3.719
co-presence	3.5	4.137	1.481
Social Dynamics	4.95	5	0.092
Avatar Appearance	3.7	3.033	-1.862
Avatar Behaviour	2.9	4.96	4.616
Other Avatars	0.631	5.35	3.660

These results show a significant improvement of the evaluation of the behaviour of the participants own avatar and of the other avatars, demonstrating that realistic animation and emotionally expressive behaviour have a strong effect on people's evaluations of avatars. There was also a significant improvement in presence and a notable but non-significant improvement in co-presence showing that this has a real effect of the participants experience. Interestingly there was a reasonably strong result that the avatars appearance was considered worse in the full animation condition. This may be because participants concentrated less on the appearance when the characters' behaviour was more lively.

Following the testing the participants were invited to make comments about their experience in an open interview. Regarding the 3D developments feedback included, "I think it's good that it's 3D because you can sort of.. it gives you more of a vision on how everyone's acting.. sort of thing and how people can react.. because you can sort of see all the shadows and stuff it's more realistic and it gets you.. into it a bit more"
Regarding the animation one pupil observed, "They did move differently, like if I said something, erm, like lovingly towards someone they did an action kind of expressing that as well"
From the teachers perspective on the critical nature of the participants, "I think the feedback's been very positive, the children are very au fait with this kind of software, they are used to using these kinds of systems at home and they can give good quality feedback on it, they know what they want, they know what they want to have and how to use it and they know what's possible."

6 Conclusion/Discussion

edrama provides a platform for participants to engage in focused discussion around emotionally charged issues. This new prototype provides an opportunity for the developers to explore how emotional issues embedded in the scenarios, characters and dialogue can be represented visually without detracting from the learning situation.
The user trials demonstrate that the creation of a 3D animated version of *edrama* indicates a marked improvement on the role-playing experience using the *edrama* system. The 3D version of the system, with the automated bit part character, may contribute to improving the perceived quality of social interaction over and above the original 2D version. In addition to this, adding emotionally appropriate animations to the user avatars improves both the participants' evaluation of those characters and their sense

of presence. There is great potential for the use of *edrama* in education in areas such as citizenship, PHSE and drama. Beyond the classroom *edrama* can be easily customised for use in professional training, where face to face training can be difficult or expensive, such as customer services training and e-learning in the workplace.
Our research shows that the application of expressive characters to online role-play contributes positively to an already engaging user experience. Future work could include the exploration of automated bit-part characters to fully develop a non-human director. Additionally tools to enable participants to replay the role-plays have been considered. These could enable further reflection and group discussion, allowing for comparisons of sessions between different groups of learners. Replays could even be altered to adjust the emotional states of each character and generate different online 'performances', which could create emotionally rich experiences for audiences as well as participants.

7 Acknowledgements

This work is a result of the LINK project eDrama - Enhancement of People, Technology and their Interaction: funded by PACCIT (People @ The Centre of Communication and Information Technologies).

The project partners would like to thank Swanshurst School, Billesley, Birmingham for their participation in this study.

8 References

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.

Briscoe, E. & Carroll, J. 2002. Robust Accurate Statistical Annotation of General Text. In *Proceedings of the 3rd International Conference on Language Resources and Evaluation*, Las Palmas, Gran Canaria. 1499-1504.

Cassell, J. & Vilhjálmsson, H. 1999. "Fully Embodied Conversational Avatars: Making Communicative Behaviors Autonomous. *Autonomous Agents and Multi-Agent Systems* 2(1): 45-64.

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*.

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.

Garau, M., Slater,M., Bee, S. and Sasse, M.A. 2001. The impact of eye gaze on communication using humanoid avatars. *Proceedings of the SIG-CHI conference on Human factors in computing systems*, March 31 - April 5, 2001, Seattle, WA USA, 309-316.

Gillies, M. & Ballin, D. 2004. Integrating autonomous behavior and user control for believable agents, in the *Proceedings of the Third international joint conference on Autonomous Agents and Multi-Agent Systems* Columbia University, New York, July 2004

Heise, D. R. 1965. Semantic Differential Profiles for 1,000 Most Frequent English Words. *Psychological Monographs*. 70 8:(Whole 601).

Mateas, M. 2002. Ph.D. Thesis. Interactive Drama, Art and Artificial Intelligence. School of Computer Science, Carnegie Mellon University.

Vilhjálmsson, H. & Cassell, J. 1998. BodyChat: Autonomous Communicative Behaviors in Avatars. Proceedings of ACM Second International Conference on Autonomous Agents, May 9-13, Minneapolis, Minnesota.

Zhang, L., Barnden, J.A., Hendley, R.J. & Wallington, A.M. 2006. Exploitation in Affect Detection in Open-ended Improvisational Text. In *Proceedings of Workshop on Sentiment and* Subjectivity at COLING-ACL 2006, Sydney, July 2006.

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.